Issue Brief

U.S. AI Statecraft

From Gulf Deals to an International Framework

Author
Pablo Chavez

CSET CENTER for SECURITY and EMERGING TECHNOLOGY

October 2025

Table of Contents

Introduction

The United States' recent artificial intelligence partnerships with the United Arab Emirates (UAE) and Saudi Arabia (KSA) represent a significant new chapter in U.S. geopolitical and economic statecraft. These multibillion-dollar agreements aim to position major U.S. technology companies—spanning cloud infrastructure, semiconductors, and advanced artificial intelligence (AI) models—at the center of both Gulf nations' AI ambitions.[1]

The KSA and UAE partnerships also serve multiple U.S. objectives. Geopolitically, they align critical digital infrastructure in the Gulf with U.S. tech instead of Chinese alternatives, reinforcing U.S. influence in a region where China has made inroads.[2] Economically, they secure new markets for advanced semiconductors and cloud services—both critical U.S. exports—while unlocking Gulf investment capital for U.S.-led AI development.[3] Strategically, the transactions potentially open a channel for the United States to influence the emerging international AI rules by embedding governance, technical safeguards, and normative alignment into the infrastructure layer of future AI ecosystems.

The Gulf partnerships are not with democracies, but they still present models for projecting democratic values into complex geopolitical environments. By incorporating transparency, accountability, security, and other mechanisms of responsible AI use into digital infrastructure, these initiatives have the potential to become durable models for U.S. AI partnerships. However, timing is critical. As China rapidly advances its own AI chip capabilities and expands its global cloud infrastructure, the window for establishing strong and strategically aligned partnerships in the Gulf and beyond may be narrowing faster than U.S. institutions and decision-making processes can adapt.[4*]

Still, if implemented effectively, the UAE and KSA partnerships could serve as prototypes for future U.S. AI engagements, laying the groundwork for responsible

[*] Throughout this issue brief, the terms "chips," "GPUs," and "semiconductors" are used interchangeably to broadly denote specialized semiconductor chips optimized for parallel computation, particularly suited to AI workloads such as training and deploying AI models.

U.S.-led AI development and governance without ceding ground to authoritarian alternatives. But success is not guaranteed. Without effective governance mechanisms and sustained U.S. leverage, these initiatives could devolve into U.S.-led AI diffusion in name only—technically American, but strategically unaligned and ultimately detrimental to U.S. interests.[5]

These concerns are particularly relevant given that the Saudi and Emirati AI developments are among the most concrete expressions to date of the Trump administration's international AI strategy, articulated in two recently released documents: America's AI Action Plan (the AI Action Plan) and the Executive Order on Promoting the Export of the American AI Technology Stack (the AI Export EO).[6] But despite emerging policy clarity, significant gaps remain, including governance frameworks, oversight mechanisms, and approaches to AI sovereignty.

Translating the Saudi and Emirati transactions and other international AI projects into a durable competitive advantage will require bridging these gaps by formalizing governance and oversight mechanisms, integrating individual deals into a coherent, principled, and scalable international framework, and coordinating AI development and deployment with allies and partners.[7]

Overview of Major U.S.–Gulf AI Initiatives (2024–2025)

The Trump administration's international AI efforts reflect a longer trajectory shaped by cooperative interests and regulatory constraints that set the stage for today's initiatives. For example, the Biden administration took important steps in the Gulf—most notably by supporting the 2024 commercial partnership between Microsoft and G42, the UAE's leading AI firm—while simultaneously introducing constraints on the international adoption of AI. Examples include the 2023 expansion of chip export controls,[8] the now-rescinded Framework for Artificial Intelligence (the AI Diffusion Rule),[9] and the launch of the voluntary Data Center Validated End-User (VEU) program, which authorizes exports of advanced computing technologies to preapproved operators under rigorous compliance measures.[10] These policies significantly shaped the regulatory environment in which current U.S.–Gulf AI partnerships emerged and have influenced the AI Action Plan and the AI Export EO.*

Against this backdrop, the recent KSA and UAE agreements have centered on building large-scale AI infrastructure, pairing Gulf funding and ambitions with U.S. technology and expertise. These countries emerged as early partners because they combine geopolitical urgency, sovereign investment capacity, and centralized governance with commercial appeal to U.S. firms, positioning them to serve as proving grounds for a new model of AI infrastructure statecraft.

These UAE and KSA AI partnerships vary significantly in their balance of U.S. operational control, governance clarity, and host-nation sovereignty over projects. Understanding their structural differences is essential to preserving the United States' influence over how these deployments are built, governed, and used.

While numerous U.S.–Gulf AI initiatives have been announced, this issue brief focuses on four transactions selected for their structural significance, the prominence of the companies involved, and geographic balance between the UAE

* The AI Action Plan and the AI Export Executive Order both articulate an international AI strategy centered on guardrails, structured compliance mechanisms, and enforceable oversight. For example, the AI Action Plan explicitly calls for robust export control enforcement, reinforcing the administration's intent to shape international AI diffusion through oversight and compliance mechanisms rather than pursue a laissez-faire approach to AI exports.

and KSA. Together, they illustrate distinct models of strategic partnership across the AI stack. The analysis draws exclusively from publicly available sources, including official policy documents, corporate announcements, and credible media reporting.[11]

Below are summaries of the transactions and their key components.

Microsoft–G42 Partnership (UAE, 2024)

In April 2024, Microsoft announced a $1.5 billion minority investment in G42,[12] gaining a board seat and deepening its strategic partnership with the UAE government-aligned AI firm.[13] As part of the agreement, G42 committed to deploy its AI applications and services on the Microsoft Azure cloud platform and to migrate key portions of its infrastructure to Azure.[14]

To address U.S. national security concerns—particularly G42's past ties to Chinese firms such as Huawei—the April 2024 partnership is governed by a "first of its kind" binding framework known as the Intergovernmental Assurance Agreement (IGAA).[15]

Negotiated with input from both the U.S. and UAE governments, the IGAA is a company-to-company agreement between Microsoft and G42 that commits the parties to comply with U.S. export control laws, "elevate" security standards, and implement safeguards against technology diversion for G42's use of Microsoft technology and the Microsoft-operated cloud and AI software environment used by G42.[16] Oversight of these commitments is exercised through a joint Microsoft-G42 compliance committee that was established through the IGAA and is meant to meet quarterly to review implementation, risk exposure, and adherence to the companies' partnership obligations.[17]

Though the IGAA text itself has not been released, public reporting identifies several guardrails tied to the deal. Microsoft has stated that, under the agreement, G42 would adopt know-your-customer protocols to exclude Chinese entities, and that it has considered technical security measures such as a vault-within-a-vault model to physically and logically isolate access to advanced AI chips and model weights.[18] In addition, the IGAA provides enforcement recourse, including financial penalties and international arbitration in London, giving Microsoft a venue for remedies outside UAE courts.[19]

G42 took additional decoupling actions both in anticipation of and in response to the IGAA.[20] These actions went beyond IGAA requirements and reflected independent (but geopolitically influenced) corporate policy decisions. For example, G42 and its affiliates committed to screening counterparties against the U.S. Consolidated Screening List (CSL), thus adding U.S.-aligned, targeted entity-specific restrictions on top of their exclusion of Chinese companies.[21] G42 also removed Huawei equipment and divested its ownership interests in various Chinese tech companies, again as voluntary corporate actions rather than pursuant to IGAA obligations.[22]

In September 2024, the U.S. government approved the export of advanced AI chips to a Microsoft-operated facility in the UAE under the Microsoft–G42 collaboration, with access and use restrictions designed to mitigate diversion risks (including China- related limits).[23] As a result, for deployments covered by the IGAA and this export approval, the operating model consists of Microsoft running the Azure cloud and AI environment inside G42 data center infrastructure**.** Where partnership workloads run in G42 facilities, they do so inside this Microsoft-operated enclave governed by the license conditions. This export approval is distinct from the larger proposed UAE AI chip exports described below.

The IGAA and the associated U.S. export approval appear to be focused only on Microsoft-operated Azure environments within G42's infrastructure. G42 AI projects that rely on non-Microsoft technology and even some deployments that use Microsoft technology but are not operated by Microsoft appear to be out of scope.[24]

Stargate UAE AI Campus (UAE, 2025)

Announced during President Trump's May 2025 visit to Abu Dhabi, Stargate UAE is a planned 10-square-mile AI supercomputing campus funded and built by G42.[25] It serves as the centerpiece of the newly launched U.S.-UAE AI Acceleration Partnership (the Acceleration Partnership), a bilateral government-to-government framework designed to facilitate technology cooperation, ensure robust U.S. security standards, and responsibly deploy advanced AI infrastructure.[26]

Other technology companies—including NVIDIA, OpenAI, Oracle, Cisco, and SoftBank—are collaborating with G42 on the first phase: a 1-gigawatt (GW) data center cluster scheduled to begin operations in 2026.[27] This initial phase is expected

to support around 100,000 of NVIDIA's next-generation Blackwell GB300 graphics processing units (GPUs), a volume of compute that would rival the scale of the world's largest AI deployments.[28] A larger 5-GW campus, if built as planned, could host at least hundreds of thousands of Blackwell GB300 chips, potentially setting new global benchmarks for AI compute capacity.

Stargate UAE is reportedly owned and funded by G42, with its physical infrastructure likely designed and built by the company's data center subsidiary, Khazna.[29] The UAE's role centers on providing land, capital, and operational hosting capacity for the project, while OpenAI and others are expected to supply and operate the AI models, software stacks, and cloud services deployed within the Stargate UAE facility. In this model, the UAE serves as the infrastructure provider and enabler, while U.S. firms deliver the core AI capabilities.[30]

U.S. officials have promoted Stargate UAE as a strategic win in securing Gulf AI infrastructure under the United States' influence rather than China's.[31]

As the first project under the Acceleration Partnership, Stargate UAE operates within this broader bilateral framework. However, unlike the Microsoft–G42 partnership, Stargate UAE lacks a formal IGAA or equivalent U.S. oversight mechanism.

Broader Context: How GPUs Flow to the UAE (2025)[32]

Under a preliminary U.S.–UAE agreement, reportedly up to 500,000 NVIDIA Blackwell-generation GPUs per year may be exported to the UAE, although the total duration and final quantity remain uncertain.

- 100,000 per year expected to go to G42-operated systems (UAE-owned).
- 400,000 per year expected to go to U.S. firms (e.g., Microsoft and Oracle) deploying infrastructure inside the UAE.
- License restrictions are expected to include: (1) prohibitions on Chinese-origin technology and personnel, and (2) requirements for logging, auditability, and U.S. inspections rights.

If fully realized, the UAE's planned GPU capacity could represent one of the world's largest concentrations of AI compute infrastructure.

Recent reporting indicates that the partnership remains under negotiation and that the export of advanced chips has not yet been authorized.[33] During Trump's May visit to Abu Dhabi, the U.S. delegation secured broad commitments from the UAE to adopt national security protocols compatible with U.S. safeguards, particularly regarding the diversion of U.S. technology. Still, detailed security conditions and enforcement mechanisms remain unresolved.[34] The United States' approval of further technology exports to Stargate UAE, and the possible adoption of more formal oversight mechanisms, will likely depend on the outcome of those ongoing negotiations.

AWS–HUMAIN AI Zone Initiative (KSA, 2025)

In May 2025, Amazon Web Services (AWS) announced a joint investment of over $5 billion with HUMAIN, a newly launched AI company owned by Saudi Arabia's Public Investment Fund (PIF) and chaired by Crown Prince Mohammed bin Salman Al Saud,[35] to build a dedicated "AI Zone"[36] as part of Saudi Vision 2030.[37] HUMAIN's stated mission is to establish KSA as a global AI leader by building a full-stack ecosystem—from infrastructure and compute platforms to specialized large language models (LLMs).[38]

The AI Zone is expected to be an additional investment beyond AWS's planned $5.3 billion Middle East (KSA) Region, a commercial hyperscale cloud region scheduled to launch in 2026.[39] The zone represents a specialized cluster focused on AI-specific workloads, equipped with advanced AI semiconductors and optimized for advanced large-scale model training, fine-tuning, and deployment.[40]

While AWS has previously deployed sovereign-compliant infrastructure—including initial services for its European Sovereign Cloud—this appears to be the company's first sovereign AI Zone explicitly structured around a national strategic partner and AI-centric design.[41] Unlike the Microsoft–G42 deal, which is governed by a formal IGAA, the AWS–HUMAIN AI Zone has no publicly disclosed company-to-company security framework in place. In addition, there is no publicly announced U.S.–KSA bilateral framework analogous to the Acceleration Partnership or any similar governance or security instrument applicable to the partnership. For the time being, oversight of technology transfer and usage appears to rely on standard U.S. export

controls, Saudi governance, and the companies' internal compliance systems.[42]

NVIDIA–HUMAIN Partnership (KSA, 2025)

HUMAIN also announced in May an agreement with NVIDIA to build a national network of AI factories as part of the Vision 2030 strategy.[43] The first phase involves the delivery of 18,000 NVIDIA Blackwell GB300 AI chips, which form the core of a planned 500-MW AI compute cluster.[44]

Saudi officials have described this deployment as a sovereign asset dedicated to training national AI models for use across government, industry, and academia.[45]

HUMAIN and NVIDIA have publicly stated plans to scale the effort to hundreds of thousands of GPUs over the next five years. More broadly, HUMAIN has stated that it aims to process 7 percent of global AI workloads by 2030, supported by a projected 6.6 GW of planned sovereign AI compute capacity.[46]

The NVIDIA-powered AI factories are a distinct infrastructure deployment from the AWS–HUMAIN AI Zone.[47] While the AI Zone will be operated by AWS using its own servers and services, the AI factories will be separately operated by HUMAIN, housing NVIDIA's chips and software stack.[48]

Broader Context: How GPUs Flow to Saudi Arabia (2025)[49]

NVIDIA has publicly announced plans to deliver an initial tranche of 18,000 NVIDIA Blackwell GB300 GPUs to HUMAIN (KSA-owned).

➡️ This initial shipment supports a planned 500-MW national AI compute campus, with HUMAIN and NVIDIA publicly targeting hundreds of thousands of GPUs over five years–projections that are ambitious but remain unconfirmed.

➡️ Future GPU exports to KSA remain subject to individual U.S. export licenses and compliance verification.

➡️ License restrictions are expected to include: (1) prohibitions on Chinese-origin technology and personnel and (2) requirements for U.S.-approved logging, auditability, and U.S. inspection rights.

The Saudi deal may currently be smaller than the UAE's, but HUMAIN aims to scale fast. With ambitions to deploy hundreds of thousands of GPUs, KSA could ultimately surpass the UAE in total AI compute.

No U.S. export license covering the planned first tranche of 18,000 GPUs has been disclosed.[50] In addition, HUMAIN is not listed as a Data Center VEU, and no bilateral IGAA-style agreement or any other governance or security instrument applicable to the NVIDIA-HUMAIN partnership (such as the Acceleration Partnership) has been announced.

Structures of U.S.–Gulf AI Partnerships

Taken together, these deals reflect a new architecture for international AI engagement: state-aligned infrastructure deployments blending U.S. technology leadership with host-nation financing and sovereign ambition. But when one looks past the announcements, these partnerships are actually quite different from one another. Each deal handles competing interests differently, balancing U.S. operational control against host-country ownership and export oversight against operational freedom. Each one also raises questions about who ultimately controls the relevant AI technology.

These partnerships represent a hybrid form of AI sovereignty. Gulf governments fund and host the infrastructure, enabling them to assert sovereign control. But in many cases, critical elements—cloud operations, access to advanced chips, and model governance—remain under the control of U.S. firms. Unlike full-stack AI sovereignty, which entails end-to-end national authority over data, infrastructure, and models, this hybrid sovereignty mixes local ownership with U.S. technology, compliance obligations, and operational influence, with elements of sovereign control layered across the stack.

The table below disaggregates these initiatives, illustrating how each combines technology, capital, and governance in a distinct way. Understanding these distinctions is essential to designing a coherent U.S. strategy for AI infrastructure statecraft.

Table 1: U.S.–Gulf AI Partnerships

Initiative	Investment Amount	Operational Control	Investor Control	U.S. Oversight	Host-Nation Control
Microsoft–G42 (UAE)	$1.5 billion minority investment by Microsoft.	The Azure cloud stack is operated by Microsoft; G42 owns and operates the data center infrastructure.	Microsoft holds a minority stake and a board seat; G42 is effectively UAE-controlled.	The Bureau of Industry and Security (BIS) oversees export licenses; the partnership is governed by an IGAA; no VEU designation has been publicly reported.	The UAE controls the physical infrastructure through G42.
Stargate UAE Campus	Multi-billion-dollar investment; exact figure undisclosed; project costs could be ~$14.4 billion (author's estimate).[51]	U.S. firms (e.g., OpenAI, Oracle) are expected to operate core AI systems; the UAE owns and operates the facilities; operational governance details are pending.	G42, which is effectively controlled by the UAE, funds the campus.	Export licenses not publicly disclosed; no IGAA; no Data Center VEU designation has been publicly reported; covered by the Acceleration Partnership (joint US-UAE governance).	The UAE is expected to control the physical infrastructure through G42.

AWS–HUMAIN AI Zone (KSA)	Over $5 billion from AWS and HUMAIN; investment split is undisclosed.	AWS will operate the cloud stack; HUMAIN will build and run LLMs; details on data center infrastructure operations have not been publicly disclosed.	Joint investment; KSA owns HUMAIN via PIF; ownership split with AWS is not publicly disclosed.	Export licenses not publicly disclosed; no IGAA; no Data Center VEU designation has been publicly reported.	AWS controls the cloud stack, but ownership of the physical infrastructure is not publicly disclosed.
NVIDIA–HUMAIN AI Factories (KSA)	No public source for the investment amount but the project cost could be up to ~$6.4 billion (author's estimate).[52]	HUMAIN operates the factories.	KSA owns HUMAIN via PIF.	Export licenses not publicly disclosed; no IGAA; no Data Center VEU designation has been publicly reported.	HUMAIN is expected to own and operate the entire compute stack.

As the comparison in Table 1 makes clear, these partnerships are not uniform in nature. Each one reflects a distinct blend of corporate control, U.S. oversight, host-nation sovereignty, and strategic risk.

The Microsoft–G42 partnership features structured governance via the IGAA, clearly delineating U.S. operational control and UAE ownership. By contrast, Stargate UAE currently relies solely on the broader Acceleration Partnership, lacking a finalized IGAA or similar binding agreement.

AWS–HUMAIN is characterized by joint investment and, to a certain extent, joint operation, with AWS expected to operate the cloud stack while HUMAIN builds and runs Saudi sovereign AI models. Unlike Microsoft-G42 and Stargate UAE, AWS–

HUMAIN has no publicly disclosed bilateral governance framework to date. Details about facility ownership and day-to-day infrastructure operations remain undisclosed, though the joint investment infrastructure suggests that the responsibility for operating the data center infrastructure might be shared rather than solely held by HUMAIN.

NVIDIA–HUMAIN represents the strongest host-nation autonomy, with HUMAIN independently operating and owning its AI facilities without a formal U.S.–UAE bilateral agreement, creating higher risks of strategic divergence.

Taken together, these diverse arrangements highlight both the opportunities and challenges of scaling U.S. influence while managing competing business interests, geopolitical frictions, host-nation sovereignty concerns, and the need for enduring governance structures, strategic coherence, and long-term trust.

What Remains Uncertain, and Why It Matters

Despite their scale and ambition, the U.S.–Gulf AI partnerships face important unresolved questions. Key governance, control, and enforcement issues remain unaddressed, creating risks that could derail their future success. What's at stake is not just implementation risk, but whether the United States will retain real influence over how its most powerful technologies are used once they leave U.S. shores.

In important respects, the AI Export EO seeks to address precisely these uncertainties. By establishing the American AI Exports program, the order moves toward a more systematic and coordinated approach to AI export governance and deployment. But even as this initiative sets out a more structured framework for promoting U.S. AI exports, it leaves many of the practical governance and implementation details undefined. As the Trump administration develops the AI Action Plan and operationalizes the AI Export EO, there is both an opportunity and a need to clarify key questions about oversight, compliance, operational control, and technological sovereignty. In particular, six critical areas of uncertainty stand out as requiring focused attention and further elaboration: (1) infrastructure control, (2) enforcement and oversight mechanisms, (3) scale-driven drift risk, (4) technology leakage and norm divergence, (5) industrial and political dependencies, and (6) fragmented governance and architectural asymmetry.

1. Infrastructure Control

In most cases, the physical infrastructure of the data centers (aside from IT hardware, including AI chips) is owned by the host country through a local partner. Who holds operational responsibility and day-to-day management of these systems varies across partnerships.

Microsoft controls its Azure-based deployment in the UAE under a formal agreement. U.S. firms are slated to manage the AI systems of Stargate UAE, though jurisdiction and compliance oversight remain unresolved. AWS and HUMAIN describe their initiative as a joint venture in which AWS will operate the cloud stack and AI services while HUMAIN builds and runs Saudi sovereign AI models, but ownership and day-to-day operations of the infrastructure have not yet been publicly defined.

In the case of the NVIDIA–HUMAIN AI factories, Saudi Arabia is positioned for full operational sovereignty from day one.

The AI Export EO's focus on structured packages presents an opportunity to explicitly address this ambiguity, though no specific framework to clarify matters such as infrastructure control or addressing sovereignty has been put forward. Export controls such as license conditions or VEU authorizations may limit how hardware is used, but they confer relatively limited ongoing control and visibility over how systems are operated or governed. Without clear legal scaffolding, host nation sovereignty over physical assets may ultimately override U.S. influence over them, regardless of the original intent.

2. Enforcement and Oversight Mechanisms

Only some of these partnerships have formal governance frameworks. The Microsoft–G42 deal is governed by the IGAA, while the Stargate UAE project falls under the Acceleration Partnership. By contrast, other deals rely on company-level compliance frameworks and (once issued) U.S. export licenses, without any required shared oversight structure or harmonized enforcement protocols across vendors.

This patchwork approach highlights the absence of standardized mechanisms for managing AI partnerships. The 2024 expansion of the VEU program was meant to standardize export compliance for trusted foreign data center operators, with provisions for preapproval, audits, and reporting. But no Gulf-based entity (or, for that matter, any entity operating outside of China and India in the case of the previously existing VEU program) appears to have been designated under the rule, perhaps reflecting a preference for bespoke agreements on strategic projects or the stringent nature of VEU requirements. The AI Export EO and the AI Action Plan open the door to more systematic compliance frameworks, though they leave key enforcement and operational oversight details largely unspecified, creating an opportunity to require VEU designations or similar standardized mechanisms for future strategic AI partnerships.

Broader questions remain about operational oversight. Even if it were used, the VEU framework does not necessarily provide the U.S. government with robust visibility

into how infrastructure is operated, how compute is allocated, or how models are trained and deployed. Furthermore, there is no effective centralized or interagency mechanism responsible for monitoring the use of AI technology across deployments outside of the United States operated or enabled by U.S. AI firms. While the AI Export EO tasks the Economic Diplomacy Action Group (EDAG) with coordinating diplomatic and financing resources for U.S. AI export projects, it does not assign EDAG or any other entity responsibility for the operational oversight or real-time compliance monitoring for exported AI technologies.

Without such centralized, coordinated, and proactive oversight, U.S. policy must rely disproportionately on reactive measures such as revoking export licenses once a significant compliance problem or strategic misalignment has already occurred. Because such actions can be politically and financially costly and occur only after damage is done, it is far preferable to establish organized oversight mechanisms that can detect and resolve issues early, preventing the need for such measures. As U.S. AI deployments abroad increase in scale and complexity, the absence of proactive, centralized oversight risks leaving U.S. policy increasingly fragmented, reactive, and ill-equipped to manage emerging challenges.

3. Scale-Driven Drift Risk

Large-scale exports of advanced AI chips may introduce a qualitatively different class of strategic risk. Even when each shipment is licensed and governed, the sheer volume of compute can overwhelm existing compliance mechanisms. With hundreds of thousands of GPUs moving abroad every year, for example, monitoring, enforcement, and audit workloads rise dramatically.[53]

Drift is the incremental, sometimes overlooked shift of a licensed deployment away from U.S. safeguards until it effectively operates beyond meaningful U.S. oversight. Consider a Gulf data center operator that starts with 100,000 licensed GPUs for a national AI factory. To meet rising demand, the operator makes a series of small capacity-enhancing changes: a few thousand additional GPUs, expanded power and cooling capacity, and one more building on the same campus. Each change is relatively minor and ostensibly compliant, yet after several rounds, the site is running well beyond its initial compute capacity without ever triggering a fresh, system-level

export review. In aggregate, the deployment has drifted into a configuration—and risk posture—never evaluated by U.S. officials.

Neither the AI Export EO nor the AI Action Plan explicitly addresses the risks or oversight challenges posed by projects of this potential magnitude. How—or even whether—the United States intends to manage drift risks tied to compute volume and the physical scale of AI campuses remains an open question. The volume and magnitude of a project may themselves constitute thresholds that warrant stand-alone governance tools, though the specific nature these scale-related risks remains uncertain.

4. Technology Leakage and Norm Divergence

Deploying U.S. AI technology abroad carries nontrivial leakage risks, meaning diversion of advanced chips, exposure of model weights, or unauthorized access through intermediaries. Governance steps can reduce, but will likely not eliminate, this risk. The UAE's de-risking efforts, including G42's severing of China ties, show serious intent to align more closely with the United States, yet reports continue to flag the possibility of diversion to Chinese companies.[54] The strategic alternative—ceding markets and partnerships to China—[55]would erode U.S. technology leadership and influence and is therefore a less preferable path, but leakage remains a structural exposure that must be assumed and managed in any overseas deployment.

Norm divergence operates as a hybrid risk that shares characteristics of both drift and leakage. It unfolds incrementally like infrastructure drift but achieves the strategic outcome of technology leakage: AI systems operating beyond meaningful U.S. influence. Once sovereign AI infrastructure is operational, U.S. visibility into how models are trained, by whom, and to what ends is inherently limited, especially in partner-operated environments. In Saudi Arabia, for example, publicly announced sovereign AI ambitions, grounded in large, locally operated facilities and plans to train national models, illustrate how host-country priorities can potentially set usage norms independently of U.S. guardrails. As national AI programs expand, they can solidify local governance practices on data use, transparency, and dual-use constraints that diverge from U.S. expectations, even if no technology ever leaks.[56]

5. Industrial and Political Dependencies

While these AI partnerships enable the U.S. to project technological power abroad, they also create a new form of dependence and potentially invert leverage. Gulf financing is helping underwrite U.S. AI companies and infrastructure, perhaps pulling them closer to Emirati and Saudi strategic priorities. If a U.S. AI chip company, for example, prioritizes sovereign deployments over U.S. public sector needs, this could affect chip availability and allocation in the United States. Meanwhile, U.S. firms might take actions they would not otherwise pursue on governance issues or human rights matters to preserve access to funding and markets. In addition, the capital and resources of foreign partners could steer U.S. policy and export decisions in directions that might not otherwise be taken. Relatedly, if host country priorities shift or funding falters, the United States could end up enmeshed in multi-billion-dollar projects it can no longer credibly oversee or influence.

6. Fragmented Governance and Architectural Asymmetry

A particularly complex risk lies in the architectural design of these partnerships. Some U.S.–Gulf partnerships, such as Microsoft–G42 and AWS–HUMAIN, embed U.S. influence directly into the cloud management layer, meaning U.S. companies manage the day-to-day operations of cloud and AI services, even though Gulf companies might own the physical data centers. Others, such as the NVIDIA–HUMAIN AI factories, transfer compute power entirely to a corporation effectively under the control of the host nation. These are not functionally equivalent. The farther the United States is from the systems where AI models are actually trained, hosted, and operated—and the more indirect its role in the systems' day-to-day operation—the harder it becomes to identify and address risk.*

* These governance differences stem primarily from varying levels of U.S. government engagement– sometimes strategic and sustained, other times reactive or absent. The Microsoft–G42 IGAA, for instance, emerged from a period of intense U.S. interagency involvement, congressional scrutiny, and national security concerns. In contrast, the AWS-HUMAIN and NVIDIA-HUMAIN arrangements indicate lighter-touch or more fragmented U.S. oversight, resulting in looser or deferred governance models. These variations underscore the need for a more consistent, policy-driven approach to AI infrastructure diplomacy.

This structural inconsistency across partnerships creates a broader governance challenge. Fragmentation is not only a product of inconsistent U.S. controls. Host nations are also asserting their sovereignty, often through legal frameworks and technical requirements that constrain how U.S. companies operate abroad. Europe is leading the way. Sovereign cloud initiatives such as France's Bleu (with Microsoft) and Germany's T-Systems' partnership with Google Cloud reflect a growing demand that U.S. infrastructure conform to local jurisdiction, data access mandates, and co-governance models.[57] These are not simply defensive. They are proactive strategies to rebalance perceived dependencies on U.S. digital technology.

As a result, U.S. companies now operate between two overlapping sovereignty regimes: Washington's evolving national security framework and host countries' sovereignty mandates. The UAE and KSA, for instance, assert their infrastructure control through financial ownership, operational governance, and co-development requirements, mirroring broader global patterns emerging across Europe and Asia.[58] Rather than a unified international model, what is developing is a patchwork of customized arrangements shaped by bilateral negotiations, sometimes resulting in architectural asymmetry characterized by a mismatch between governance frameworks and AI deployments.

The AI Action Plan and the AI Export EO emphasize the strategic value of exporting comprehensive U.S. AI technology packages and embedding U.S. standards into an international AI infrastructure, but both documents leave unanswered how the U.S. government will manage tensions when exported AI systems must simultaneously comply with U.S. governance requirements and host nations' sovereignty demands.

This architectural disconnect is exacerbated by a broader lack of clarity around U.S. AI decision-making. It is sometimes unclear why certain countries are subject to tighter restrictions or what conditions would allow for looser ones. In some cases, risk assessments appear to be based on a country's overall relationship with the United States or its connections to adversaries like China, while in others, the intended use of the AI systems—such as the potential for dual use in military, surveillance, or intelligence contexts—raises flags.

These distinctions are rarely explained publicly. For countries that are given restricted

access to high-performance compute, for example, it has not always been evident whether the barrier is political, technical, or precautionary. This was a central criticism of the now-defunct AI Diffusion Rule. Key partners, including strategic allies, expressed concern that they were excluded without a clear articulation of whether the barrier was political alignment, technical risk, or use-case sensitivity.[59]

Unless the United States leverages the opportunity created by the AI Action Plan and the AI Export EO to clearly articulate its AI decision-making criteria—and develops mechanisms to transparently differentiate between, for example, use-case risks and country-specific risks—it will continue to face ambiguity in its own decision-making processes. This ambiguity risks alienating allies, strengthening adversaries, and weakening the effectiveness of the United States' international AI strategy. Until the United States can formalize a model that accounts for these distinctions, it will continue to generate friction in its international AI partnerships.

Policy Recommendations: A U.S. Framework for an International AI Infrastructure

To secure the strategic gains of U.S.-Gulf AI partnerships—and responsibly extend the model to other regions—the United States must move from fragmented, ad hoc deal-making to a principled, durable framework for AI infrastructure cooperation. This requires moving beyond case-by-case negotiations to establish standardized frameworks that protect U.S. interests while respecting partner sovereignty.

At its core, this strategy should be built around a clear system for determining who gains access to advanced AI technologies and under what conditions. The goal is not rigid classification, but a flexible, risk-calibrated approach that can transparently differentiate among partners, address evolving threats, and extend across regions and architectures.

The United States already has pieces of this framework: export controls, licensing regimes, bilateral agreements, and the Data Center VEU program. But it lacks an overarching strategy to integrate these tools into a coherent model of governance. The recommendations that follow aim to provide that structure.* These recommendations would help operationalize the administration's stated approach, articulated in the AI Action Plan and the AI Export EO, of exporting packages that potentially span multiple layers of the AI technology stack, thereby embedding governance and oversight more effectively.

1. Establish a Structured, Rules-Based Framework for Access to U.S. AI

While the AI Action Plan introduces procedural structure through the U.S. AI Exports Program, the United States should go further by establishing transparent, rules-based criteria determining which countries gain access to advanced chips, AI compute infrastructure, and frontier models.

---

* These recommendations do not seek to address all the identified risks. Instead, they propose governance mechanisms adaptable to emerging challenges as partnerships evolve. The objective is not to codify the bespoke arrangements seen in current U.S.–Gulf AI partnerships, but to incorporate their most effective safeguards into a durable, flexible framework for future international AI partnerships.

This framework should rest on four core pillars: (1) strategic alignment, to ensure geopolitical coherence, mutual interest, and long-term stability of partnerships that advance U.S. interests; (2) intended use, to distinguish beneficial civilian and commercial applications from those posing potentially unacceptable risks to U.S. interests, such as military or mass-surveillance uses; (3) institutional capacity, to ensure safeguards that protect U.S. technology and investments; and (4) deployment scale, recognizing that large-scale AI infrastructure deployments—measured by factors such as power capacity, compute density, and operational complexity—may overwhelm traditional compliance mechanisms, necessitating specialized monitoring and governance approaches.

Moreover, the framework must balance thoroughness with speed. To implement it or analogous structures, the United States will need to leverage both its expertise and resources and its willingness to streamline decision-making processes to outpace competitors. If the United States does not make and implement AI governance decisions quickly and efficiently, it risks losing out to Chinese firms that continue to deploy AI infrastructure abroad. Huawei, Alibaba, Baidu, and others do not wait for U.S. decision-making, and slow and difficult-to-implement governance mechanisms risk ceding advantage to competitors offering simpler, less-constrained alternatives.

Countries with close defense and intelligence ties to the United States—particularly allies with strong rule-of-law institutions and proven compliance records—should continue to qualify for deeper and faster integration. For these trusted partners, access to advanced compute should be the default setting, provided that baseline safeguards are in place. These might include, for example, the exclusion of Chinese hardware, restrictions on diverting hardware or services to countries of concern, and civilian-use attestations.

Institutionally capable partners, in particular, offer the United States practical advantages—including predictability in compliance and adherence to the rule of law— that protect U.S. investments and broader U.S. interests. By working preferentially with proven partners, the United States effectively mitigates risks by leveraging the strengths of transparent and accountable governments.

By contrast, countries lacking institutional capacity, credible oversight mechanisms, or strategic alignment with the United States should not receive presumptive access. Access in these cases should be narrower, more conditional, and governed by externally enforced safeguards that protect U.S. interests. Technical requirements—such as localized telemetry or third-party audits—may resemble those required of allies in certain circumstances, but the verification and enforcement burdens should fall more heavily on the United States to protect its technological advantages. Where host-country legal systems cannot credibly support oversight, access should be conditioned on direct contractual control, run-time monitoring, and limited exposure to frontier capabilities. If core risks to U.S. technology and security cannot be mitigated, access should be denied until solutions are found.

This access framework should function as a long-term strategic instrument, not just as a one-time export screening tool. It should provide reliable allies and partners with a transparent path into U.S.-aligned AI ecosystems, while reinforcing governance norms at the infrastructure layer. Rather than viewing safeguards as bureaucratic obstacles, the framework would treat them as shared responsibilities in a secure, strategically aligned system that benefits both U.S. interests and those of allies and partners.

The proposed approach builds on the stated objective of the now-rescinded AI Diffusion Rule—managing the proliferation of high-risk AI capabilities—while addressing criticisms that the rule's design was rigid and overbroad and emphasizing the importance of promoting U.S. AI exports. Where the Diffusion Rule imposed inflexible national categories, this model enables calibrated access based on actual risk and demonstrated trustworthiness, as well as geopolitical value to the United States. Done well, such a framework would balance predictability with geopolitical flexibility, thus creating guardrails, not roadblocks. It would also help rebuild trust among partners who have criticized recent U.S. export controls as arbitrary or ad hoc. A structured, transparent model would potentially help reestablish the United States as a reliable (and not just powerful) technology partner.

Access must be conditional, but it should also be principled and trust-based. A robust democracy using AI to provide health services should not face the same

restrictions as an autocracy building predictive surveillance systems. If the United States fails to distinguish among potential partners and embed those distinctions in its governance framework, it risks alienating allies, eroding normative leadership, and empowering actors who could repurpose U.S. AI in ways that ultimately threaten U.S. interests.

2. Build a Layered Governance Architecture for AI Infrastructure

A credible access framework must be coupled with enforceable and sustainable safeguards. The United States should explore establishing a layered governance architecture for AI infrastructure exports that combines strategic oversight, regulatory consistency, and operational enforcement. This architecture could consist of three distinct but interrelated levels:

(1) Strategic Alignment via Government-to-Government (G2G) Agreements. G2G agreements define the outer perimeter of cooperation, establishing when and how the United States will authorize AI infrastructure partnerships with particular countries. The Acceleration Partnership provides an early example of this G2G approach, though its details are not fully disclosed and may remain under development. These agreements could include restrictions on geographic deployment, compute volume thresholds, and national security safeguards. This layer would anchor bilateral engagement in shared strategic objectives and trust, drawing on precedents such as defense and security cooperation agreements that establish foundational terms and conditions between governments.*

(2) Regulatory Baselines Through Standardized Compliance Mechanisms. Tools such as the Data Center VEU authorization provide a relatively standardized and rules-based pathway for exports, especially to higher-risk or strategically complex jurisdictions. The existing Data Center VEU framework is essentially a pre-clearance list for trusted data center infrastructure outside of the United States, enabling predictable export conditions while enhancing transparency and potentially mutual

* While G2G agreements would require diplomatic engagement, the framework envisioned here assumes a standardized baseline agreement that could be applied across multiple countries. This would reduce the need for fully bespoke negotiations and ensure consistency, while allowing for country-specific provisions where needed.

trust. If a data center is on the list, it can receive approved AI hardware under a set of consistent rules, including audit and inspection rights for the U.S. government and reporting obligations about how the data center in question is used. Currently, Data Center VEU participation is voluntary. However, for large-scale deployments or transactions involving partners for which governance risks are elevated, the United States should consider making VEU participation a baseline requirement and, more broadly, the administration should consider reducing the VEU program's compliance complexity. By doing so, the United States would ensure consistent safeguards, reduce regulatory ambiguity and burdens, and establish a clearer foundation for trust even with partners whose more limited strategic alignment or institutional capacity necessitate more stringent oversight.

(3) Operational Enforcement Using Company-Level Agreements. Company-level agreements—such as Microsoft's IGAA with G42—introduce enforceable obligations when a U.S. firm is directly engaged in a commercial or operational partnership with a host-country entity.[*] In such cases, these contracts can embed export conditions into day-to-day operations by specifying matters such as operational access controls and shared governance structures. They offer a flexible means of managing risk via contractual commitments. Companies operating in sensitive sectors often use contracts to embed detailed compliance obligations on their partners to mitigate risks. In this context, the value of these contractual mechanisms is in part their integration into a structured, layered governance architecture.

Together, these three governance layers would form a scalable governance system. G2G agreements set strategic boundaries, including deployment scale thresholds, geographic limitations, and usage restrictions. Regulatory mechanisms such as the Data Center VEU provide a consistent compliance infrastructure across multiple deployments. And company-level contracts operationalize those commitments in specific partnerships, embedding safeguards into contractual terms and technical controls. Such a layered architecture offers the United States a flexible yet systematic approach capable of scaling across different AI infrastructure deployments and

[*] The Microsoft–G42 partnership illustrates an early—though largely untested—version of this layered approach, combining government-facilitated negotiations, the IGAA, and export licensing conditions.

partnership configurations.

Critically, this governance architecture must prioritize speed of implementation. Perfect frameworks that arrive after competitors have seized market share serve no strategic purpose. The objective is to establish workable partnerships that can mature and expand over time. Such an approach transforms export controls from simple regulatory mechanisms into comprehensive policy instruments that serve diplomatic, economic, and security goals concurrently.

The modularity of this governance architecture is particularly valuable given the varying degrees of host-nation sovereignty in emerging AI partnerships. Arrangements in which U.S. companies maintain direct operational oversight might rely more heavily on company-level agreements and lean less on G2G agreements and other mechanisms. By contrast, deployments that are characterized by greater host-country control and minimal direct U.S. operational presence may require increased reliance on robust G2G frameworks and more stringent regulatory baselines to ensure compliance and strategic alignment.

The modular approach also provides flexibility over time. For example, in the case of trusted partners, it allows for tailored options such as the expansion of the existing Data Center VEU program into jointly managed bilateral frameworks. Rather than relying solely on unilateral U.S. oversight, bilateral VEUs would enable close allies to actively participate in matters such as defining security standards and sharing enforcement responsibilities.[*]

This cooperative approach aligns closely with current digital sovereignty initiatives among European allies, enhancing regulatory coherence and potentially reducing bilateral friction. Though diplomatically and technically complex, bilateral VEUs (perhaps beginning with pilot programs) would concretely signal mutual trust and offer a compelling diplomatic alternative to perceptions of unilateral U.S. gatekeeping while still providing a robust foundation for safeguarding sensitive AI

[*] For example, even though differences remain, there is potential alignment between U.S. and European cybersecurity frameworks in areas such as facility security standards. Exploring and capitalizing on these overlaps could facilitate bilateral cooperation and regulatory coherence.

infrastructure.

As AI infrastructure becomes increasingly hybrid in ownership and control, sovereignty is emerging not as binary but as layered. One actor may own the physical facilities, another may operate the cloud stack, and a third may govern access to AI models. Such division of control increasingly requires a governance approach that is both comprehensive and responsive to these market dynamics. Without embedded and enforceable safeguards, U.S. influence risks diminishing precisely when sovereign AI ambitions worldwide are intensifying.

3. Align Institutional Capacity with Strategic Objectives

Even the most carefully developed governance model is unlikely to succeed without a government capable of implementing it. Although broader resource and budget considerations are deeply important, addressing critical personnel gaps is particularly urgent. Agencies responsible for export licensing, compliance enforcement, economic diplomacy, and partnership development currently lack sufficient specialized staffing and global operational reach.[60] These shortfalls create bureaucratic friction, exacerbate compliance vulnerabilities, and impede geopolitically important deals precisely when speed and operational effectiveness matter most.

To close these gaps and effectively pursue the strategic objectives outlined in the AI Action Plan and the AI Export EO—business development, risk mitigation, and diplomatic influence—the United States should: (1) recruit specialized personnel at the Departments of State and Commerce, Office of Science and Technology Policy, and the EDAG skilled in areas such as economic diplomacy, public-private partnerships, and AI sector-specific knowledge; (2) expand Bureau of Industry and Security (BIS) staffing and technical teams dedicated specifically to AI licensing, compliance enforcement, infrastructure oversight, and technical evaluation of hardware and software; (3) deploy dedicated AI-focused officers at key U.S. embassies for real-time compliance monitoring, proactive diplomatic engagement, and rapid enforcement responsiveness; and (4) appoint experienced diplomatic and technical personnel to lead international coordination efforts revolving around AI deployment and governance.

These investments are strategic rather than merely administrative. Without the

capacity of government personnel to effectively manage international AI deployments, even the most principled policy frameworks will erode under the weight of complexity, scale, and geopolitical pressure.[61] Equally important is implementation speed. Bureaucratic delays risk losing potential partners to Chinese alternatives that might impose fewer governance requirements.

4. Launch an AI Cooperation Forum to Build Strategic Alignment

To move beyond case-by-case agreements and promote long-term alignment, the United States should lead the creation of an AI Cooperation Forum (AICF), a flexible international platform designed to support trusted countries in developing sovereign AI capabilities while embedding enforceable governance safeguards.*

The AICF would prioritize rapid deployment of collaborative projects, recognizing that speed is essential to prevent partners from defaulting to faster-moving Chinese AI infrastructure options. The forum would coordinate public-private investment in national AI projects, share best practices on secure infrastructure deployment, and link bilateral partnerships to broader international AI development and governance initiatives, all while implementing basic safeguards first and adding governance sophistication over time.

The AICF would not impose a one-size-fits-all model. Instead, countries could develop AI frameworks that connect with others while fitting their own capabilities, strategic needs, and political realities. This sovereignty-respecting approach would draw on lessons from both successful national strategies and early-stage bilateral deals. By serving as a connective framework between those efforts, the AICF would help align trusted partners around shared principles while reinforcing U.S. leadership in shaping a secure, rules-based international AI ecosystem.[62]

* The Gulf approach may not work elsewhere. Gulf countries can deploy massive sovereign wealth funds and make rapid decisions through centralized leadership, enabling them to commit billions to AI partnerships within months. By contrast, democratic governments face legislative approval processes, budget constraints, and multi-stakeholder decision-making that can delay or limit such commitments. Future analysis should consider how different countries' institutional structures and resource capabilities affect their ability to participate in large-scale AI partnerships.

Table 2 summarizes how the key strategic risks identified in this analysis align with the specific policy recommendations.

Table 2: Mapping Key Risks to Policy Recommendations

Risk Category	Strategic Concern	Policy Options
Infrastructure Control	U.S. firms may nominally operate AI infrastructure abroad but lack clearly defined jurisdiction over infrastructure, data, and operations.	Clearly define operational control, legal jurisdiction, and ongoing compliance oversight in G2G agreements, company-level agreements (such as an IGAA), and standardized tools such as the VEU program.
Enforcement and Oversight Mechanisms	Compliance currently relies heavily on company-level frameworks and standard export licenses without centralized oversight or harmonized enforcement protocols across vendors.	Adopt a layered governance architecture combining G2G agreements, standardized regulatory baselines, enforceable company-level agreements, and expanded institutional capacity at BIS, EDAG, and other agencies.
Scale-Driven Drift Risk	Incremental expansions of AI deployments can cumulatively shift projects beyond original oversight parameters, undermining U.S. safeguards and strategic alignment.	Ensure the structured, rules-based framework addresses volume- and scale-related risks by incorporating mechanisms to reassess cumulative expansions and potentially require enhanced oversight safeguards.

Technology Leakage and Norm Divergence	Host could divert chips or models to third parties of concern or repurpose U.S.-exported AI systems in ways that fall outside U.S. visibility and interests.	Require explicit use-case attestations, operational safeguards, and embedded auditability within structured frameworks, standardized regulatory mechanisms, and company-level agreements to minimize leakage and norm divergence.
Industrial and Political Dependencies	Host-country variables like financing and chip demand may distort U.S. priorities and actions, constrain U.S. leverage, or invert U.S. influence.	Promote transparency and explicitly align public-private strategic priorities through oversight mechanisms embedded within licensing frameworks and contractual arrangements.
Fragmented Governance and Architectural Asymmetry	Bespoke deals create regulatory fragmentation and architectural asymmetry (misaligning control of infrastructure with governance) that reduce U.S. visibility; opaque export criteria alienate allies and create unpredictability.	Establish the AICF to coordinate public-private AI investments and align governance standards; pair it with standardized decision criteria and baseline regulatory tools to reduce fragmentation and asymmetry across deployments.

Conclusion: From Transactions to Frameworks

The U.S.–Gulf AI partnerships mark a new phase of compute-enabled diplomacy. Geopolitically, these deals help displace Chinese digital infrastructure from critical systems, embedding U.S. technology and influence at the core of emerging AI ecosystems. Economically, they secure markets for essential U.S. exports and unlock significant Gulf capital for U.S.-led AI development.

More subtly, they function as programmable soft power. By extending access to compute and embedding U.S. platforms into international digital ecosystems, the U.S. hard-wires long-term interdependence, projecting influence not just through presence but also through infrastructure. If foreign governments, companies, and other organizations adopt and rely on the U.S. AI stack, the United States increases its chances of achieving enduring technological advantage, economic benefits, and normative influence.

To achieve these goals, the United States must build on the AI Action Plan and AI Export Executive Order and shift from transactional deals to a principled framework for AI-infrastructure statecraft. That framework should specify eligibility for advanced compute, baseline safeguards, and enforceable oversight mechanisms that reconcile sovereignty concerns systematically. While today's deal-driven approach has yielded quick wins, its transactional nature offers limited guardrails against shifting geopolitical circumstances. Given the stakes, success should be judged not by any single deal or political moment, but by whether the United States builds a transparent, enforceable architecture that endures across administrations.

A structured architecture that combines strategic agreements, standardized regulatory tools, and company-level safeguards offers a scalable and modular foundation for AI export governance. If done right, this model could become the backbone of a secure and interoperable international AI ecosystem anchored in shared values, resilient to misuse, and aligned with U.S. strategic interests. But without clear rules and credible safeguards that are implemented swiftly, even the most ambitious partnerships will struggle to endure, and the United States might thus ultimately cede technological and normative leadership to less trustworthy actors.

The Gulf partnerships offer critical test cases. How the United States manages these relationships and applies the lessons learned will determine whether American AI leadership translates into lasting strategic advantage or becomes a cautionary tale of missed opportunities.

Author

Pablo Chavez is a non-resident senior fellow at CSET. He is a technology policy expert focused on artificial intelligence and cloud computing. Chavez serves on the board of the Open Technology Fund, is an adjunct senior fellow at the Center for a New American Security, and advises corporations and nonprofits on technology policy and strategy. Previously, he led global public policy at Google Cloud, U.S. public policy at Microsoft, and global public policy and government affairs at LinkedIn. Earlier in his career, Chavez served as a senior staffer in the U.S. Senate. He holds a law degree from Stanford Law School and a bachelor's degree in public and international affairs from Princeton University.

Acknowledgments

The author would like to express his sincere gratitude to everyone who contributed to the completion of this research. The resources and support provided by CSET and the CSET team were especially instrumental in making this work possible. Special thanks go to Emelia Probasco, Owen Daniels, Jacob Feldgoise, Shelton Fitch, Jason Ly, and Chloe Moffett for their invaluable contributions. The research and analysis in this report also benefited significantly from the policy expertise and thought leadership of the team at the Center for a New American Security, for which the author is deeply grateful.

Author Disclosure

The views and opinions expressed in this essay are solely those of the author and do not necessarily reflect the views of any organizations or institutions with which the author works or is affiliated.

Document Identifier: doi: 10.51593/20250019

Endnotes

[1] Alison Snyder and Dave Lawler, "Trump's Gulf Gamble: Helping UAE and Saudi Become AI Powers," *Axios*, May 18, 2025, https://www.axios.com/2025/05/18/trump-gulf-ai-deals-saudi-uae-security-china-risk.

[2] Will Knight, "The Middle East Has Entered the AI Group Chat," *Wired*, May 15, 2025, https://www.wired.com/story/trump-middle-east-artificial-intelligence-investments; Huawei, "Huawei Cloud in Saudi Arabia," https://activity.huaweicloud.com/intl/en-us/saudi_arabia_region.html.

[3] Semiconductor Industry Association, "SIA 2024 Factbook (2024)," https://www.semiconductors.org/wp-content/uploads/2024/05/SIA-2024-Factbook.pdf; Council of Economic Advisers, "What Drives the U.S. Services Trade Surplus? Growth in Digitally Enabled Services Exports," The White House, June 10, 2024, https://bidenwhitehouse.archives.gov/cea/written-materials/2024/06/10/what-drives-the-u-s-services-trade-surplus-growth-in-digitally-enabled-services-exports.

[4] Huawei is reportedly making rapid progress with its Ascend 910C chips. See Anton Shilov, "DeepSeek Research Suggests Huawei's Ascend 910C Delivers 60% of Nvidia H100 Inference Performance," *Tom's Hardware*, February 4, 2025, https://www.tomshardware.com/tech-industry/artificial-intelligence/deepseek-research-suggests-huaweis-ascend-910c-delivers-60-percent-nvidia-h100-inference-performance. Other Chinese firms also continue advancing AI chip development despite U.S. export controls. See Kyle Chan, Gregory Smith, Jimmy Goodrich et al., "Full Stack: China's Evolving Industrial Policy for AI," (RAND, June 26, 2025), https://www.rand.org/pubs/perspectives/PEA4012-1.html. Meanwhile, Chinese cloud companies are expanding across Southeast Asia and other parts of the world. See Debby Wu, "Alibaba Expands AI Cloud Services in Malaysia, Philippines," *Bloomberg*, July 1, 2025, https://www.bloomberg.com/news/articles/2025-07-02/alibaba-expands-ai-cloud-services-in-malaysia-philippines.

[5] President Trump recently announced at a press conference that export licenses for NVIDIA's H20 and AMD's MI308 AI chips to China would be approved on the condition that the companies pay a 15% revenue-share to the U.S. government. While this arrangement would likely generate revenue for the United States government, it exemplifies the risk of "pay-to-export" models: the dilution of U.S. strategic leverage and normalization of advanced technology transfers with limited governance oversight. See White House, "President Trump Holds a Press Conference, Aug. 11, 2025," video, August 11, 2025, https://www.whitehouse.gov/videos/president-trump-holds-a-press-conference-aug-11-2025; Associated Press, "US will get a 15% cut of Nvidia and AMD chip sales to China under a new, unusual agreement," August 11, 2025, https://apnews.com/article/nvidia-

amd-15-revenue- share-deal-c06e20d9c3418f1d0b1292891c4610c6.

[6] Executive Office of the President, *America's AI Action Plan*, (Washington, DC: The White House, July 23, 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf; Exec. Order No. 14320, 90 FR 35393 (2025), https://www.whitehouse.gov/presidential-actions/2025/07/promoting-the-export-of-the-american-ai-technology-stack.

[7] The Trump administration's strategic intent behind exporting U.S. AI technology was recently articulated by Michael Kratsios, director for the White House Office of Science and Technology Policy, who stated, "We want the entire world to be running on an American artificial intelligence stack…our cloud, our chips, our algorithm, all of that needs to be exported and packaged to the world, so that we become the ecosystem of choice globally." Stephanie Lai and Edward Ludlow, "Trump's AI Plan Seeks to Have US Set Global Standard, Aides Say," *Bloomberg*, July 24, 2025, https://www.bloomberg.com/news/articles/2025-07- 24/trump-s-ai-plan-seeks-to-have-us-set-global-standard-aides-say. While this statement emphasizes strategic alignment, it also highlights the necessity for governance frameworks and compliance mechanisms, as discussed in this issue brief.

[8] "New Export Controls on Advanced Computing and Semiconductor Manufacturing: Five Key Takeaways," Sidley Austin LLP, November 1, 2023, https://www.sidley.com/en/insights/newsupdates/2023/10/new-export-controls-on-advanced-computing-and-semiconductor-manufacturing.

[9] Bureau of Industry and Security, "Department of Commerce Announces Recission [sic] of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls," U.S. Department of Commerce, May 13, 2025, https://www.bis.gov/press-release/department-commerce-announces-recission-biden-era-artificial-intelligence-diffusion-rule-strengthens-chip. For the full text of the rule, see "Framework for Artificial Intelligence Diffusion," 90 FR 4544 (January 15, 2025), https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion. Note that the Bureau of Industry and Security (BIS) has *announced* the rescission of the AI Diffusion Rule and directed non-enforcement pending a replacement rule. However, until a rescission notice is formally published in the *Federal Register*, the rule technically remains in the Code of Federal Regulations even though it is not being applied in practice.

[10] Bureau of Industry and Security, "Expansion of Validated End User Authorization: Data Center Validated End User Authorization," 89 FR 80080 (October 2, 2024), https://www.govinfo.gov/content/pkg/FR-2024-10-02/pdf/2024-22587.pdf. For an explanation of the Data Center VEU authorization, see Paul, Weiss, Rifkind, Wharton & Garrison LLP, "U.S. Government Eases Export Control Restrictions for AI Chips Bound for Qualifying Data Centers," October 3, 2024, https://www.paulweiss.com/insights/client-memos/us-government-eases-export-control-restrictions-for-ai-chips-bound-for-qualifying-data-centers. For operational details, see 15

C.F.R. § 748.15 and Supplements Nos. 7, 8, and 9. As of mid-2025, none of the Gulf initiatives described in this issue brief appears to rely on the Data Center VEU authorization.

[11] While one of the major transactions examined in this brief—the Microsoft-G42 agreement—was finalized during the Biden administration in 2024, it represents the type of U.S.-Gulf AI partnership that the Trump administration has prioritized and offers valuable structural and governance lessons for current policy approaches.

[12] "Microsoft Invests $1.5 Billion in Abu Dhabi's G42 to Accelerate AI Development and Global Expansion," Microsoft, April 16, 2024, https://news.microsoft.com/source/2024/04/16/microsoft-invests-1-5-billion-in-abu-dhabis-g42-to-accelerate-ai-development-and-global-expansion;

[13] Sheikh Tahnoon bin Zayed Al Nahyan—the UAE's National Security Advisor, the deputy ruler of Abu Dhabi, and the brother of the UAE's president—is the majority shareholder of Group 42 Holding Ltd. (the full name of G42), indicating controlling ownership of the entity. Mubadala Investment Company (Abu Dhabi's sovereign wealth fund) also holds a minority stake, further reinforcing UAE state-linked control. Additional minority investors include Silver Lake, a U.S.-based private equity firm. U.S. Securities and Exchange Commission, Schedule 13D (vTv Therapeutics), filed June 10, 2022, https://www.sec.gov/Archives/edgar/data/1641489/000110465922070134/tm2218155d2_sc13d.htm; Mubadala, "Mubadala takes stake in Group 42," November 2, 2020, https://www.mubadala.com/en/news/mubadala-takes-stake-group-42; Silver Lake, "G42 announces investment by Silver Lake," April 14, 2021, https://www.silverlake.com/g42-announces-investment-by-silver-lake. The Microsoft board seat is held by Brad Smith, Microsoft's Vice Chair and President. "Tahnoon bin Zayed Al Nahyan Chairs G42's Final Board," Emirates News Agency-WAM, November 7, 2024, https://www.wam.ae/en/article/b62qzsb-tahnoon-bin-zayed-nahyan-chairs-g42%E2%80%99s-final-board.

[14] "Microsoft and G42 Partner to Accelerate AI Innovation in UAE and Beyond," Microsoft (blog), April 15, 2024, https://blogs.microsoft.com/blog/2024/04/15/microsoft-and-g42-partner-to-accelerate-ai-innovation-in-uae-and-beyond.

[15] Microsoft, "Microsoft and G42 Partner to Accelerate AI Innovation in the UAE and Beyond." The IGAA governs G42's use of Microsoft technology but does not appear to extend to G42's broader operations or infrastructure. G42 continues to design, build, and operate independent AI infrastructure through subsidiaries such as Core42 and Khazna. See Jason Ma, "Khazna to Build Data Center Near Ankara, Turkey," *Data Center Dynamics*, April 30, 2025, https://www.datacenterdynamics.com/en/news/khazna-to-build-data-center-near-ankara-turkey. Although G42's commitment to remove Huawei equipment and sever ties with Chinese tech suppliers appears to extend to its subsidiaries and its investments, the precise breadth and depth of G42's

decoupling from China remains unclear. See Omar El Chmouri, "Nvidia Teams Up with G42 for UAE Data Centers in Mideast AI Push," *Bloomberg*, June 11, 2025, https://www.bloomberg.com/news/articles/2025-06-11/nvidia-teams-up-with-g42-for-uae-data-centers-in-mideast-ai-push; Chloe Cornish and Kaye Wiggins, "Abu Dhabi AI Group G42 Sells Its China Stakes to Appease US," *Financial Times*, February 9, 2024, https://www.ft.com/content/82473ec4-fa7a-43f2-897c-ceb9b10ffd7a. Moreover, other actors in the UAE reportedly continue to pursue AI partnerships with Chinese firms that fall outside the scope of U.S. oversight agreements and may continue to pose strategic and governance challenges. See "OpenAI Says China's Zhipu AI Gaining Ground amid Beijing's Global AI Push," *Reuters*, June 25, 2025, https://www.reuters.com/world/china/openai-says-chinas-zhipu-ai-gaining-ground-amid-beijings-global-ai-push-2025-06-25. These nuances underscore the need for careful governance of these AI partnerships and deployments.

[16] Microsoft, "Microsoft and G42 Partner to Accelerate AI Innovation in UAE and Beyond."

[17] Microsoft, "Microsoft and G42 Partner to Accelerate AI Innovation in UAE and Beyond."

[18] Stephen Nellis, "Exclusive: Microsoft's UAE deal could transfer key U.S. chips and AI technology abroad," *Reuters*, May 24, 2024, https://www.reuters.com/world/middle-east/microsofts-uae-deal-could-transfer-key-us-chips-ai-technology-abroad-2024-05-23.

[19] In response to congressional oversight, Microsoft also restructured elements of its commercial partnership with G42 to tighten U.S. control, specifically by shifting to leasing rather than transferring certain sensitive AI products. U.S. House of Representatives, Select Committee on the Strategic Competition Between the United States and the Chinese Communist Party, "Letter to NSA on Microsoft's Billion-Dollar Partnership with UAE Firm G42," July 11, 2024, https://selectcommitteeontheccp.house.gov/media/letters/letter-nsa-microsofts-billion-dollar-partnership-uae-firm-g42; John Sakellariadis, "Microsoft Scales Back AI Partnership with Emirati Firm amid Concerns over China Ties," *Politico*, August 1, 2024, https://www.politico.com/news/2024/08/01/microsoft-ai-partnership-emirati-china-concern-00172167.

[20] See Trevor Hunnicutt and Alexandra Alper, "Microsoft–G42 Deal Positive Because It Cut Huawei Ties—White House Official," *Reuters*, June 24, 2024, https://www.reuters.com/technology/microsoft-g42-deal-positive-because-it-cut-huawei-ties-white-house-official-says-2024-06-24.

[21] "Another Important Step in Advancing Responsible AI to Serve the World," Microsoft, September 17, 2024, https://news.microsoft.com/source/2024/09/17/another-important-step-in-advancing-responsible-ai-to-serve-the-world. The CSL is a U.S. government list identifying parties that are subject to federal restrictions on the export, reexport, or transfer of goods and items. See U.S. Department of Commerce, International Trade Administration, "*Consolidated*

Screening List, https://www.trade.gov/consolidated-screening-list.

[22] Cornish and Wiggins, "Abu Dhabi AI Group G42 Sells Its China Stakes to Appease US."

[23] Alison Snyder and Maria Curi, "Scoop: Advanced AI Chips Cleared for Export to UAE Under Microsoft Deal," *Axios*, December 7, 2024, https://www.axios.com/2024/12/07/us-uae-microsoft-g42-ai-chips.

[24] For example, G42 continues to design, build, and operate independent AI infrastructure through subsidiaries such as Core42 and Khazna. See Jason Ma, "Khazna to Build Data Center Near Ankara, Turkey," *Data Center Dynamics*, April 30, 2025. It is also partnering with other technology companies to build AI data centers. See Omar El Chmouri, "Nvidia Teams Up with G42 for UAE Data Centers in Mideast AI Push," *Bloomberg*, June 11, 2025, https://www.bloomberg.com/news/articles/2025-06-11/nvidia-teams-up-with-g42-for-uae-data-centers-in-mideast-ai-push. G42 is also partnering with Microsoft on cloud and AI projects that appear to fall outside the scope of the IGAA. See Georgia Butler, "Microsoft and Core42 to build sovereign cloud and AI infrastructure for Abu Dhabi," *Data Center Dynamics*, March 20, 2025, https://www.datacenterdynamics.com/en/news/microsoft-and-core42-to-build-sovereign-cloud-and-ai-infrastructure-for-abu-dhabi.

[25] Stephen Nellis, "'Stargate UAE' AI Datacenter to Begin Operation in 2026," Reuters, May 22, 2025, https://www.reuters.com/business/media-telecom/stargate-uae-ai-datacenter-begin-operation-2026-2025-05-22.

[26] U.S. Department of Commerce, "UAE/US Framework on Advanced Technology Cooperation," May 15, 2025, https://www.commerce.gov/news/press-releases/2025/05/uae/us-framework-advanced-technology-cooperation; UAE Ministry of Foreign Affairs, "UAE-US: A Strategic Partnership Built on Five Decades of Mutual Trust," May 17, 2025, https://www.mofa.gov.ae/en/MediaHub/News/2025/5/17/17-5-2025-UAE-us; El Chmouri, "Nvidia Teams Up with G42 for UAE Data Centers."

[27] Eliot Brown, "OpenAI Commits to Giant U.A.E. Data Center in Global Expansion," *Wall Street Journal*, May 22, 2025, https://www.wsj.com/tech/open-ai-abu-dhabi-data-center-1c3e384d; Nellis, "'Stargate UAE' AI Datacenter to Begin Operation in 2026."

[28] Analyst estimates indicate that NVIDIA's GB300 NVL72 rack-scale system, which includes 72 GB300 GPUs, draws approximately 140 kilowatts (kW) per rack. Based on this figure, just 200 MW of capacity (roughly 1,400 racks) could support around 100,000 GPUs, which implies that Stargate UAE's planned 1 GW buildout in Phase 1 could accommodate five times that amount (about 500,000 GPUs) or more. See "NVIDIA GB300 to Feature Enhanced Specifications, Full Rack

Shipments Expected to Gradually Scale in 3Q25, Says Trendforce," TrendForce, March 18, 2025, https://www.trendforce.com/presscenter/news/20250318-12522.html; Federico Maccioni, Manya Saini, and Yousef Saba, "UAE to Build Biggest AI Campus Outside US in Trump Deal, Bypassing Past China Worries," Reuters, May 15, 2025, https://www.reuters.com/world/china/uae-set-deepen-ai-links-with-united-states-after-past-curbs-over-china-2025-05-15.

[29] Matthew Tostevin, "Building AI's Backbone in the Middle East," *Newsweek*, July 24, 2025, https://www.newsweek.com/nw-ai/building-ais-backbone-middle-east-2102446; Khazna Data Centers, "Facilities," accessed July 12, 2025, https://www.khazna.ae/facilities. See also Kelsey Warner, "UAE's Biggest Data Center Firm Eyes US Expansion," *Semafor*, May 2, 2025, https://www.semafor.com/article/05/02/2025/uaes-biggest-data-center-firm-eyes-us-expansion.

[30] Primary jurisdiction over data centers on UAE soil remains with the UAE unless explicitly altered through formal bilateral agreements granting limited jurisdictional rights to the United States. Absent such agreements, U.S. oversight relies solely on contractual and export license conditions. See Alexander Cornwell, "US–UAE Multi-Billion Dollar AI Data Campus Deal Far from Finalized, Sources Say," *Reuters*, June 6, 2025, https://www.reuters.com/business/finance/us-uae-multi-billion-dollar-ai-data-campus-deal-far-finalised-sources-say-2025-06-06.

[31] David Sacks (@davidsacks47), X, May 17, 2025, https://x.com/davidsacks47/status/1923716659497283709. Mr. Sacks currently serves as the White House AI and Crypto Czar. He also leads the President's Council of Advisors on Science and Technology (PCAST).

[32] Karen Freifeld and Hadeel Al Sayegh, "US Close to Letting UAE Import Millions of Nvidia's AI Chips, Sources Say," Reuters, May 14, 2025, https://www.reuters.com/business/finance/us-close-letting-uae-import-millions-nvidias-ai-chips-sources-say-2025-05-14; Brown, "OpenAI Commits to Giant U.A.E. Data Center."

[33] Amrith Ramkumar and Eliot Brown, "National-Security Concerns Tie Up Trump's U.A.E. Chips Deal," *Wall Street Journal*, July 16, 2025, https://www.wsj.com/politics/national-security/national-security-concerns-tie-up-trumps-u-a-e-chips-deal-a0273815.

[34] Cornwell, "US-UAE Multi-Billion Dollar AI Data Campus Deal"; Ramkumar and Brown, "National-Security Concerns Tie Up Trump's U.A.E. Chips Deal."

[35] Ahmed Al Omran and Andrew England, "Saudi Arabia launches AI venture Humain ahead of Donald Trump visit," *Financial Times*, May 12, 2025, https://www.ft.com/content/2082b1e5-e571-42a5-96b6-e7d5c2977afa.

[36] "AWS and HUMAIN Announce a More Than $5B Investment to Accelerate AI Adoption in Saudi

Arabia and Globally," Amazon, May 13, 2025, https://www.aboutamazon.com/news/company-news/amazon-aws-humain-ai-investment-in-saudi-arabia.

[37] Public Investment Fund, "HRH Crown Prince Launches HUMAIN as Global AI Powerhouse," May 12, 2025, https://www.pif.gov.sa/en/news-and-insights/press-releases/2025/hrh-crown-prince-launches-humain-as-global-ai-powerhouse.

[38] HUMAIN homepage, accessed July 12, 2025, https://www.humain.ai/en.

[39] Amazon, "AWS and HUMAIN Announce a More Than $5B Investment."

[40] Amazon, "AWS and HUMAIN Announce a More Than $5B Investment."

[41] Max Peterson, "Announcing Initial Services Available in the AWS European Sovereign Cloud, Backed by the Full Power of AWS," AWS Security Blog, July 3, 2024, https://aws.amazon.com/blogs/security/announcing-initial-services-available-in-the-aws-european-sovereign-cloud-backed-by-the-full-power-of-aws. AWS is building a second AI Zone in South Korea in partnership with SK Group. See SK Group, "SK Group and AWS Team Up to Build Cloud Computing Infrastructure to Support AI Innovation," June 21, 2025, https://eng.sk.com/news/sk-group-and-aws-team-up-to-build-cloud-computing-infrastructure-to-support-ai-innovation.

[42] Amazon, "AWS and HUMAIN Announce a More Than $5B Investment." U.S. oversight primarily rests on conditions stipulated at export approval, though BIS retains legal authority to conduct post-shipment verifications, audits, and periodic compliance checks. However, given practical resource constraints, these may occur infrequently relative to export volumes. BIS, housed within the Department of Commerce, is responsible for implementing and enforcing export controls on dual-use items, including advanced semiconductors and AI-related technologies. BIS regulates exports under the Export Administration Regulations, issues export licenses, and oversees compliance through mechanisms such as the VEU program. See Bureau of Industry and Security, U.S. Department of Commerce, https://www.bis.doc.gov.

[43] Max A. Cherney and Stephen Nellis, "US Tech Firms Nvidia, AMD Secure AI Deals as Trump Tours Gulf States," *Reuters*, May 14, 2025, https://www.reuters.com/world/middle-east/saudi-arabia-partners-with-nvidia-spur-ai-goals-trump-visits-2025-05-13.

[44] Kif Leswing, "Nvidia Sending 18,000 of Its Top AI Chips to Saudi Arabia," *CNBC*, May 13, 2025, https://www.cnbc.com/2025/05/13/nvidia-blackwell-ai-chips-saudi-arabia.html.

[45] Cherney and Nellis, "US Tech Firms Nvidia, AMD Secure AI Deals."

[46] Chloe Cornish, Melissa Heikkilä, and Ahmed Al Omran, "Can the Gulf Really Become an AI Superpower?" *Financial Times*, June 1, 2025, https://www.ft.com/content/509e1b95-9fe9-4402-b97a-7c2c9ba9a2f6; Cherney and Nellis, "US Tech Firms Nvidia, AMD Secure AI Deals"; NVIDIA, "HUMAIN and NVIDIA Announce Strategic Partnership to Build AI Factories of the Future in Saudi

Arabia," NVIDIA Newsroom, May 14, 2025, https://nvidianews.nvidia.com/news/humain-and-nvidia-announce-strategic-partnership-to-build-ai-factories-of-the-future-in-saudi-arabia; Dylan Patel, Jeremie Eliahou Ontiveros, AJ Kourabi, Ivan Chiam, and Wega Chu, "AI Arrives in the Middle East: US Strikes a Deal with UAE and KSA," SemiAnalysis, May 16, 2025, semianalysis.com/2025/05/16/ai-arrives-in-the-middle-east-us-strikes-a-deal-with-uae-and-ksa.

[47] NVIDIA, "Saudi Arabia and NVIDIA to Build AI Factories to Power Next Wave of Intelligence for the Age of Reasoning," NVIDIA Newsroom, May 13, 2025, https://nvidianews.nvidia.com/news/saudi-arabia-and-nvidia-to-build-ai-factories-to-power-next-wave-of-intelligence-for-the-age-of-reasoning.

[48] Aaron Raj, "HUMAIN Taps NVIDIA, AMD, AWS, Cisco and Qualcomm to Build Saudi Arabia's AI Infrastructure," *CRN Asia*, May 14, 2025, https://www.crnasia.com/news/2025/artificial-intelligence/humain-taps-nvidia.

[49] Cherney and Nellis, "US Tech Firms Nvidia, AMD Secure AI Deals." In parallel, HUMAIN and AMD announced a $10 billion joint venture to build additional AI infrastructure based on AMD hardware. "AMD and HUMAIN Form Strategic $10B Collaboration to Advance Global AI," AMD, May 14, 2025, https://ir.amd.com/news-events/press-releases/detail/1250/amd-and-humain-form-strategic-10b-collaboration-to-advance-global-ai.

[50] Christine Burke, "Saudi's Humain to Open Data Centers With US Chips in 2026," *Bloomberg*, August 25, 2025, https://www.bloomberg.com/news/articles/2025-08-25/saudi-s-humain-to-open-data-centers-with-us-chips-in-early-2026.

[51] While no official estimate has been released, press reports describe the Stargate UAE campus as a "multi-billion dollar" initiative. Cornwell, "US-UAE Multi-Billion Dollar AI Data Campus Deal." The project plans to deploy 100,000 NVIDIA GB300 chips at an estimated $51,000 to $56,000 each based on recent system-level pricing for NVIDIA's GB300 NVL72 systems, totaling $5.1 to $5.6 billion for core hardware (actual unit costs may vary significantly based on customer, volume, and other variables). Georgia Butler, "Apple to spend $1bn on Nvidia GB300 NVL72 systems: report," *Data Center Dynamics*, March 31, 2025, https://www.datacenterdynamics.com/en/news/apple-to-spend-1bn-on-nvidia-gb300-nvl72-systems-report. Infrastructure costs for the planned 1-GW facility can be estimated at approximately $8.8 billion for infrastructure, bringing total project costs to $13.9 to $14.4 billion (the index shows UAE construction costs at $8.80 per watt in 2024). Turner & Townsend, "Data Centre Cost Index 2024," October 2024, https://reports.turnerandtownsend.com/dcci-2024. This estimate represents a significant cost advantage compared to global averages and helps explain the similarity to G42's reported investment of $8 to $10 billion, which may cover the initial phase or represent a partial investment in the full project. See Dan Primack, "Top of the Morning," *Axios Pro Rata*, May 22, 2025,

https://www.axios.com/newsletters/axios-pro-rata-fd49f847-b35f-49ba-b78d-a50316e9b51a.

[52] The 18,000 GB300 AI chips alone would cost between $925 million and $1 billion, based on recent system-level pricing for NVIDIA's GB300 NVL72 systems. Butler, "Apple to Spend $1bn on Nvidia GB300 NVL72 Systems." Using Turner & Townsend's 2024 data center cost index, which shows KSA construction costs at $10.80 per watt, the 500-MW facility would cost approximately $5.4 billion in infrastructure alone, bringing total project costs to $6.3 to $6.4 billion. Turner & Townsend, "Data Centre Cost Index 2024." This initial 500 MW is significantly more than needed for 18,000 GPUs (using the 140 kW per rack estimate referenced above in note 28), suggesting the facility is designed for substantial future expansion.

[53] This observation builds on analysis by Center for New American Security (CNAS) researchers and policy experts who argue that restricting access to advanced AI chips—particularly in high-volume deployments—is essential to mitigating proliferation and ensuring responsible global AI development. See Vivek Chilukuri, Michael Depp, Bill Drexel, Janet Egan, Paul Scharre, Josh Walin, Becca Wasser, and Caleb Withers, "Response to Request for Information on the Development of an Artificial Intelligence (AI) Action Plan" (Center for a New American Security, March 2025), https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/CNAS-AI-Action-Plan-Submission.pdf.

[54] Andrew G. Clemmensen, Rebecca Redlich, and Grant Rumley, "G42 and the China-UAE-U.S. Triangle" (Washington Institute for Near East Policy, April 3, 2024), https://www.washingtoninstitute.org/policy-analysis/g42-and-china-uae-us-triangle.

[55] Tripp Mickle and Ana Swanson, "Outsourcer in Chief: Is Trump Trading Away America's Tech Future?" *New York Times*, May 15, 2025, https://www.nytimes.com/2025/05/15/business/economy/trump-chips-ai-uae.html.

[56] See Alasdair Phillips-Robins and Sam Winter-Levy, "How to Gauge Whether Trump's AI Chip Deals with Gulf Countries Are Any Good" (Carnegie Endowment for International Peace, May 16, 2025), https://carnegieendowment.org/emissary/2025/05/ai-chip-trump-gulf-uae-saudi-security-risk-good-deal.

[57] Georgia Butler, "Google Reaffirms Sovereign Cloud Solutions for EU: Aims to Ease Anxieties About Trade Tensions," *Data Center Dynamics*, May 22, 2025, https://www.datacenterdynamics.com/en/news/google-reaffirms-sovereign-cloud-solutions-for-eu. For a deeper analysis of European approaches to digital sovereignty and cloud governance, see Pablo Chavez, "European Digital Sovereignty vs. U.S. Cloud Power," *Consensus Drift*, June 3, 2025, https://consensusdrift.substack.com/p/european-digital-sovereignty-us-cloud-409.

[58] Recent AI transactions in Asia include Google Cloud's 2024 partnership with Malaysia's DNeX to deploy locally governed, air-gapped cloud services, and Nvidia's 2024 strategic alliance with India's Reliance Industries to build domestically controlled AI supercomputing infrastructure. See Danial Azhar, "DNeX, Google Cloud Sign Multi-Year Deal to Provide Sovereign Cloud Services in Malaysia," *Reuters*, September 30, 2024, https://www.reuters.com/technology/artificial-intelligence/dnex- google-cloud-sign-multi-year-deal-provide-sovereign-cloud-services-malaysia-2024-09-30/; and Munsif Vengattil, "Nvidia Partners with India's Reliance, Tata to Develop AI Applications," *Reuters*, September 8, 2023, https://www.reuters.com/technology/nvidia-reliance-partners-develop-ai-apps-india-2023-09-08.

[59] See Pieter Haeck, "Poland Fumes as US Blocks AI Chip Exports," *Politico*, January 21, 2025, https://www.politico.eu/article/poland-fumes-us-block-joe-biden-ai-chips-cap-export.

[60] For example, recent testimony before Congress has emphasized the need for increased funding for BIS to implement and enforce AI export controls. Witnesses have highlighted shortfalls in staffing and global enforcement capacity and supported BIS's FY 2025 budget request for $223.4 million. See Samuel Hammond, "Supporting the Bureau of Industry and Security" (Foundation for American Innovation, May 12, 2024), https://www.thefai.org/posts/supporting-the-bureau-of-industry-and-security; Marc Selinger, "BIS Nominee Says He Would Review New AI Chip Controls, Agency Funding," *Export Compliance Daily*, February 28, 2025, https://exportcompliancedaily.com/article/2025/02/28/bis-nominee-says-he-would-review-new-ai-chip-controls-agency-funding-2502270041.

[61] A 2008 report by the U.S. Government Accountability Office (GAO) identified significant capacity shortfalls in BIS's implementation of the VEU program, including the lack of a G2G agreement with China for on-site reviews and the absence of standardized procedures for conducting inspections. GAO recommended halting the program until BIS could ensure sufficient resources and oversight mechanisms were in place. U.S. Government Accountability Office, "Export Controls: Challenges with Commerce's Validated End-User Program May Limit Its Ability to Ensure That Semiconductor Equipment Exported to China Is Used as Intended," GAO-08-1095 (Washington, D.C.: U.S. Government Accountability Office, September 25, 2008; publicly released October 27, 2008), https://www.gao.gov/assets/gao-08-1095.pdf.

[62] This version of the AICF builds on an earlier proposal for a sovereignty-respecting, values-aligned forum for international AI cooperation. Pablo Chavez, "Sovereign AI in a Hybrid World: National Strategies and Policy Responses," *Lawfare*, November 7, 2024, https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses. On July 26, 2025, at the World Artificial Intelligence Conference in Shanghai, Premier Li Qiang proposed the creation of a global AI cooperation organization and unveiled a *Global AI Governance Action Plan*. The proposal is similar to the AICF concept. Both envision a multilateral

platform to coordinate AI development and governance. Brenda Goh, "China proposes new global AI cooperation organisation," *Reuters*, July 26, 2025, https://www.reuters.com/world/china/china-proposes-new-global-ai-cooperation-organisation-2025-07-26/; Ministry of Foreign Affairs of the People's Republic of China, "Global Artificial Intelligence Governance Action Plan," July 26, 2025, https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202507/t20250729_11679232.html.