Redefining National Security Governance through Access and Compute: U.S.-South Korea Export Controls on AI Infrastructure

By ChangHee Kim

AI has emerged as a core strategic asset, and its performance depends less on the models themselves than on the data centers and high-performance compute resources—the total computational power, measured in graphics processing unit (GPU) hours, used to train and operate AI systems—that support them. Such infrastructure can enable the diffusion of technology through mere access, without any physical transfer of equipment, creating a structural challenge for export control regimes designed on the premise of physical exports.

The U.S. and South Korean export control systems remain structured around the cross-border movement of tangible items and technical documents. In a borderless AI infrastructure environment, however, the actual route of technology transfer is determined less by where servers are located than by who can access which compute resources and data, and with what entitlements. Current export control regimes do not adequately capture this reality.

South Korean companies making large-scale investments in AI infrastructure in the United States are simultaneously subject to U.S. export rules, which treat the release of controlled technology to foreign persons within the United States as an export, and to South Korea's strategic technology control system. In this process, the same behavior—for example, a South Korean engineer remotely accessing a U.S. data center—can be interpreted differently under the two systems, amplifying regulatory uncertainty and investment risk.

This paper argues for a transition from hardware-centric export controls to compute-centric export governance. It proposes redefining AI data centers not as assets located inside or outside physical borders, but as "borderless strategic assets" managed through access entitlement structures. It discusses how Zero Trust—a cybersecurity principle that requires identity verification and authorization at each access point—and real-time access control can apply to

---

ChangHee Kim is a global trade compliance professional with over 20 years of experience specializing in U.S. and South Korean export controls (ITAR/EAR and Korean strategic items regulations), reexport compliance, and national security governance.

The author is grateful to Attorney Sejin Jung of Lee & Ko for helpful comments on legal and policy issues related to U.S. export controls and South Korean strategic technology regulations. Attorney Jung is a leading expert in Korean digital and data law, with particular expertise in the AI Basic Act and data governance legislation, and has contributed to shaping Korea's AI governance discourse through his published work in the field. Any remaining errors are the author's own.

national security governance. Specifically, the paper: 1) shows that technology transfers in AI data centers are shifting from physical movement to access entitlement structures; 2) identifies a growing misalignment between data protection law and export control law; 3) proposes compute quota and risk-tiered Zero Trust-based access governance as new policy control concepts; and 4) links functional weaponization of AI to export and access control debates.

Recent policy developments in both countries emphasize the importance of these policy questions. In March 2026, the U.S. Department of Commerce drafted rules that would require licenses for virtually all AI chip exports globally, adding conditions to transfers of more than 200,000 chips on recipient countries agreeing to build AI data centers in the United States.[1] In South Korea, the AI Basic Act entered into force on January 22, 2026, establishing risk management obligations for high-performance AI systems, with phased enforcement beginning in 2027.[2] Seoul also began distributing the first tranche of a 10,000-unit national GPU pool in March 2026.[3] Taken together, these developments create new opportunities for policy coordination between the United States and South Korea, but also expose governance gaps that neither side is yet well-positioned to manage alone.

The End of Physical Borders and the Rise of Compute

The problem of AI governance fundamentally clashes with assumptions about technology transfer on which traditional export control systems were built. Those systems evolved on the premise that strategic technologies physically cross borders. From the Coordinating Committee for Multilateral Export Controls—the informal multilateral export control regime of the Cold War—to the post-Cold War Wassenaar Arrangement, multilateral regimes have focused on controlling the movement of physical equipment and technical documents.[4] Semiconductors, precision machinery, and telecommunications equipment remain among the key controlled items today.

This challenge has taken on new urgency under the current policy environment. In January 2025, the Donald Trump administration issued an executive order prioritizing U.S. leadership in AI infrastructure, signaling a continued emphasis on compute access within U.S. national security policy.[5] In South Korea, the Lee Jae Myung administration has similarly identified AI and semiconductor supply chains as strategic priorities, pursuing domestic AI capability development while deepening investment ties with U.S. cloud and data center ecosystems.[6] These parallel policy imperatives—U.S. efforts to maintain technological leadership and South Korea's ambitions to secure AI infrastructure capacity—create both alignment opportunities and governance gaps.

The spread of AI technologies is now substantively shaking the physical-transfer premise of export control regimes. Today, strategic AI capabilities can rapidly diffuse without physical transfers, through access to compute resources and data. For example, if an actor has long-term access to a high-performance GPU cluster—an array of specialized processors optimized for the parallel computations required to train and run AI models—and continuously uses it to train and refine models, it can accumulate strategic AI capabilities without any movement of hardware.

In this environment, the central axis of national security risk is shifting from the question of "What has moved?" to "Who can access which compute resources?" Consider the case of 100 identical GPUs. If they are shared among research institutions in multiple allied countries, rather than being monopolized by a single organization in one location, the physical configuration may be the same, but the security implications can be very different. The real risk is determined by how much compute a particular actor can concentrate and for how long. (For a more detailed explanation, see Appendix A.)

As the AI market grows, data centers and high-performance compute infrastructure have become core elements of strategic competition. South Korean firms have become major investors and operators in the U.S. AI ecosystem, yet current export control systems were not designed to address compute access as a new channel for technology diffusion.[7]

This paper is informed by broader compliance and policy debates in export control and AI governance. Policy discussions and compliance commentary increasingly highlight deep uncertainty in cloud-based AI infrastructure environments that cannot be easily addressed using traditional metrics such as server location or equipment counts, with growing concern that legitimate technical activities may fall into a regulatory gray zone.[8] The purpose of this paper is therefore not to offer a final institutional blueprint but to identify the policy questions that now require closer U.S.-South Korea coordination. In particular, it seeks to explore a new governance framework that allows security and industrial innovation to coexist by focusing on two axes: access rights and compute quotas.

The paper makes three main points. First, AI infrastructure is no longer transferred primarily through the movement of hardware, but through access. High-end AI capabilities can spread across borders solely via remote access to data center resources and compute—without the physical relocation of equipment.

Second, simply allowing long-term, repeated use of compute resources can enable the transfer and accumulation of strategic capabilities, even in the absence of any formal provision or disclosure of the underlying technology. Even with identical hardware configurations, an actor's cumulative use of compute over time can yield a qualitatively different level of functional AI capability.

Third, in response to these changes, export control regimes should move beyond a perspective centered on physical equipment and technical documentation, and reconsider computation and access as core units of control. The argument here is not that existing systems should be replaced, but that they no longer fully capture the security risks and compliance uncertainties posed by AI infrastructure environments.

In practice, the current export regime appears markedly less effective at capturing computation- and access-based risks in cloud and multi-tenant environments—where different organizations share the same physical infrastructure but are logically separated through virtualization and access controls—than in traditional on-premises settings.

AI Data Centers and the Borderless Nature of Access-Based Transfers

In AI infrastructure environments, technology transfer can no longer be adequately described in terms of physical movement or the transmission of documents alone. In large-scale AI data centers and cloud-based computing infrastructures, it is not the physical location of servers but rather who can access which compute resources and datasets, and with what kind of authorization, that effectively determines the scope of the "transfer."

In the case of multinational enterprises, data centers operated by the U.S. affiliate of a South Korean organization may be located either in South Korea or in the United States, and the nationality and residence of personnel with access rights to those data centers are highly diverse.[9] In these settings, the actual pathways through which technology and data move are defined by access-control structures rather than national borders, and usage patterns and associated risks are difficult to understand based solely on physical location.

The Prior Evolution of Data Protection Law

Data protection law in some countries has already incorporated this reality to a significant extent. Both the European Union's General Data Protection Regulation (GDPR) and South Korea's Personal Information Protection Act (PIPA) interpret cross-border transfers of personal data as encompassing not only the physical transmission of data, but also conduct that makes data accessible from abroad.[10] Even if servers are located domestically, supervisory authorities generally apply regulations governing cross-border transfers if overseas personnel can remotely view or process personal data.[11]

This regulatory logic is not confined to data protection. It is gradually being extended to other regulatory regimes that govern the transfer of technology and industrial capabilities.[12]

Moreover, under the U.S. export control system, this kind of access-based use is discussed as a core issue in determining how technology should be controlled.[13] For example, consider a scenario in which an AI model developed and trained on a server in Seoul is remotely accessed by an affiliate researcher in Silicon Valley via a virtual private network (VPN) to analyze its architecture or conduct additional training. Even if the AI model and associated know-how are not physically transferred abroad, there is a question of whether export control rules treat this as a transfer of technology if foreign personnel can, in practice, access and utilize that technology. Ultimately, in AI infrastructure environments, the criteria for technology transfer are shifting from the storage location of technology or the routing of data flows toward who can access and exploit that technology.

Structural Limitations of Export Control Regimes

By contrast, the U.S. Export Administration Regulations (EAR) and International Traffic in Arms Regulations (ITAR), as well as South Korea's systems for controlling strategic goods and strategic technologies, remain grounded in a traditional concept of technology provision and disclosure and do not fully reflect today's reality.[14] The U.S. export regime similarly treats the

"release" of controlled technology to foreign persons in the United States as an export, but its regulatory focus lies on whether the substantive contents of the technology or software have been disclosed.[15]

As a result, there are interpretive gaps regarding how to distinguish and regulate access to high-performance GPU resources, on the one hand, and the transfer of controlled technical information, on the other. These ambiguities leave room for regulatory arbitrage and create burdens on both policymakers and industry by undermining companies' ability to predict what levels of access and collaboration are permitted.

South Korea's regime is also framed primarily around active conduct, such as the provision, transmission, and teaching of technology, and does not explicitly capture access-based usage patterns as a distinct regulatory unit. For example, whether the act of a South Korean engineer utilizing GPUs in a U.S. data center via Application Programming Interface (API) calls—standardized requests sent over a network to access and trigger specific functions in a remote system—constitutes the provision of a controlled strategic technology is, under the current text of the law, open to interpretive dispute.

Furthermore, both traditional forms of remote access, such as encrypted Secure Shell (SSH) connections or console logins, and API calls that expose high-risk functions can amount to substantive access to compute in AI infrastructure environments.[16] For instance, if a South Korean engineer can repeatedly invoke APIs such as "/train," "/fine-tune," or "/target-detection" against a model deployed in a U.S. data center to perform military- or security-related AI functions, this could result in a significant level of functional capability being transferred to and concentrated in a particular actor, even in the absence of any transfer of technical documentation. Nevertheless, current export control rules do not explicitly treat such API-based access to computation as an independent regulatory category, creating a gray area that will require dedicated discussion to design U.S.-South Korea joint governance mechanisms.[17]

Consequently, while an access-rights-based notion of transfer has been substantially codified in the field of data protection, the export control space remains anchored in a physical and act-based framework, creating potential regulatory misalignment. This misalignment should be addressed through close policy and governance coordination within the U.S.-South Korea alliance.

Compute Quotas: A New Policy Control Concept and Its Institutional Implications

Current export control regimes focus on acts of providing or disclosing technology and on the physical movement of equipment. The actual amount of compute used is not yet treated as an independent regulatory unit. Existing debates have mainly emphasized the number of GPUs and the performance of individual devices, but in practice, AI capability is determined by the cumulative compute quota—the total amount of AI compute resources that a particular actor uses over a specified period—consumed over time.[18]

Here, compute quota refers to the combination of usage time, the number of devices used in parallel, and workload intensity that yields functional capability. It is this capability that must become the object of control if the United States and South Korea are to manage substantive technology transfer and capability accumulation in a borderless infrastructure environment.

For example, assume that an identical configuration of 100 H100 GPUs is available. If one organization runs them for only two days on an experimental basis, while another organization uses the same configuration almost every day for six months to train large-scale models, the visible hardware is the same, but the magnitude and character of the AI capability accumulated are completely different. The former is likely to remain at the level of short performance tests or demos, whereas the latter can repeatedly train and refine massive models, rapidly building strategic-level AI capabilities that can be directly deployed for military and security missions such as precise target identification, situational awareness, and cyber operations automation.

This paper treats compute quota not simply as a technical measure, but as a policy concept that may help define the upper bound of permitted AI capability accumulation. Rather than controlling individual equipment movements, the idea is to manage the upper bound of AI functionality that a particular actor may accumulate. Instead of focusing solely on controlling physical transfers, the key question for export licensing and joint governance centers on how much AI compute will be permitted. Once compute quota is introduced as a concept, policymakers can move beyond asking about the destination and quantity of GPU exports and focus instead on how much countries, organizations, and projects use AI compute. This can serve as a useful policy tool for preventing specific countries or actors from concentrating and accumulating strategic AI capabilities over long periods.

At the same time, compute alone does not determine the quality or significance of a technology transfer. In practice, strategic AI capability also depends on model architecture, data quality and availability, engineering expertise, organizational learning, and the operational context in which systems are deployed. Compute quota should therefore not be understood as a universal proxy for all forms of AI risk, but as a policy handle that is most informative where sustained access to high-performance compute is coupled with frontier model training, capability concentration, and high-risk functional deployment. In other domains, such as narrow, data-limited applications or environments where talent and organizational capacity are the binding constraints, compute-centric controls should be supplemented by other regulatory instruments and risk indicators.

Institutional Gaps in U.S.-South Korea Export Controls

The export control systems of the United States and South Korea increasingly recognize that high-performance computing resources are a key determinant of AI capability. However, neither country has yet to institutionalize compute quotas as an independent control concept. The United States indirectly controls compute capacity through export restrictions on high-end GPUs and data center equipment, but regulations focus on performance thresholds and quantity limits for individual hardware items.[19] South Korea's system likewise centers on the

provision and transfer of technology, making it difficult to systematically capture the cumulative and continuous use of compute.[20]

As a result, if an actor maintains the same hardware configuration but conducts large-scale cumulative compute over an extended period, it can accumulate strategic AI capabilities within a regulatory gray zone. Even if GPUs are imported under an export license and then used to provide compute to a specific actor that far exceeds what was initially anticipated, current regimes have limited ability to detect or manage that change.

The compute quota concept addresses these gray areas and offers a starting point for U.S.-South Korea joint governance discussions. Concretely, this would mean shifting the regulatory focus from whether equipment has been brought in to who is continuously using how much compute, and incorporating compute quotas as a core variable in licensing, notification, and reporting frameworks. For allied countries, a tiered structure could be designed in which higher ceilings and more flexible operation are permitted, creating an institutional basis for the United States and South Korea to jointly monitor and manage the concentration of sensitive compute resources.

Borderless Strategic Assets and Zero Trust as a Policy Logic

Current export control systems rely primarily on *ex ante* regulation, through document submissions and licensing reviews, combined with *ex post* reporting and audits. This approach can be effective for controlling imports and exports of physical equipment, but it faces clear limitations in cloud-based AI infrastructure, where access to compute can shift rapidly. Once a license is approved, it is difficult to track and adjust, in real time, who is actually using how much compute under which conditions.

In the United States and Europe, regulators have already begun tightening controls on high-end AI chips and related technical data, while introducing requirements for companies to systematically retain and manage access logs and usage records for such data.[21] U.S. rules still focus heavily on whether technology has been provided or disclosed and whether license conditions have been complied with, so logs are commonly used for *ex post* audits and compliance checks.[22] Deemed export rules and interpretations on cloud and remote access likewise hinge on whether technical data has been "released," and therefore do not fully capture newer risks, such as API-based compute access or large-scale concentration of compute.[23]

The U.S. AI Diffusion Rule reflected some awareness of these access- and compute-based risks. The framework's conditions on transfers of advanced computing hardware included measures such as Validated End-User (VEU) status, clustering limitations, ongoing security and logging obligations, and restrictions on certain forms of model training and infrastructure use for non-allied destinations. These measures implicitly address aspects of AI infrastructure by tying hardware exports to monitoring, acceptable-use controls, and access management for a limited group of trusted countries, including South Korea. However, they remain fundamentally anchored in hardware export licensing and do not yet treat access-permission structures

and cumulative compute quotas as primary units of control in their own right. The approach proposed in this paper is therefore not a rejection of the framework, but an attempt to extend its logic by making access and compute explicit, first-order variables in U.S.-South Korea joint governance of AI infrastructure.

This paper draws on the Zero Trust concept from cybersecurity and applies it as a policy framework for export control and national security governance.[24] Zero Trust is the principle of not presuming that any person or organization is trustworthy, but instead verifying identity and authorization at each point of access and granting only the minimum privileges necessary for the task at hand. In this paper, Zero Trust serves as a conceptual framework for articulating where access should be allowed and where it should be restricted in a borderless AI environment.

Zero Trust's most stringent forms—continuous verification, fine-grained attribute-based policies, real-time enforcement, and strict compute ceilings—are particularly appropriate where access involves sensitive models, high-end training functions, defense-related applications, or high-risk APIs. By contrast, lower-risk collaborative research among trusted institutions in allied countries may warrant lighter application of the same principles, with greater reliance on *ex post* auditing and institutional safeguards. A risk-tiered application of Zero Trust thus offers a more realistic path for embedding access-centric controls into existing export control systems without imposing disproportionate burdens on benign innovation.

This perspective makes it possible to consider real-time access control as a new policy option, alongside existing document-based licensing. Incorporating real-time access control does not mean abolishing the existing licensing and reporting framework, but rather embedding the principles agreed at the licensing stage directly into data center access control and logging systems. This could involve: 1) assigning attributes such as nationality, affiliation, and project purpose to each account and role; 2) defining policies based on those attributes—for example, "Accounts of specified nationalities may not access training functions on high-end GPU clusters," or "Allied public research project accounts may use up to X GPU-hours per month"; and 3) recording and monitoring all access and compute usage in real time, so that access is automatically blocked and alerts triggered if predefined thresholds are exceeded. In such a system, principles agreed upon by regulators and companies are automatically enforced in data center operations.

As AI data centers become more sophisticated, U.S.-South Korea export control systems must be redesigned so that access entitlement structures, rather than physical borders, become the core unit of control. Building on existing document- and license-based systems, a step-by-step integration of Zero Trust and real-time access control offers a realistic and implementable policy pathway for dealing with AI infrastructure.

Expanding Export Control Norms to Reflect AI Functional Weaponization Risks

AI weaponization is not confined to software embedded in physical weapon systems such as autonomous lethal weapons. From the moment AI begins performing a range of military and security functions—such as target identification, situational awareness, decision support, and cyber operations automation—it functionally acquires a weapon-like character.[25]

Of course, physical weapon platforms such as autonomous unmanned aerial vehicles (UAVs) and ballistic missile systems have long been treated as core controlled items under existing export control regimes, including ITAR and EAR.[26] What this paper emphasizes, however, is that actors can use essential weapon functions such as target identification, route planning, and situational awareness solely through access to high-performance compute and sensitive data, negating the need for systems embedded in hardware. At this point, compute quotas and access-permission structures emerge as a new axis for controlling weaponization.

This kind of functional weaponization is possible even without transferring the source code or the model itself, relying instead only on access to compute resources and the operating environment. For example, if an actor merely secures access to an AI analysis system linked to sensitive sensor data, it can use that system to analyze and predict an adversary's military activities.[27] This is difficult to fully capture within the traditional export control regime, which is structured around the transfer of physical items and technical documentation.

Compute access, model deployment locations, and the degree of integration with military and intelligence systems must be examined as new control points. In other words, where a given model is deployed and who can use it should be treated as core variables in assessing functional weaponization risk. By doing so, debates on AI weaponization need not remain at the abstract level of ethics and norms; instead, the concepts of compute quota and access-permission structures can be directly connected to the concrete design of export-control and access-control mechanisms.

Conclusion and Policy Questions

As technologies and compute resources diffuse across borders, a U.S.-South Korea joint export control system—policy alignment plus joint governance—serves as a realistic alternative to effective AI export control.

From a compliance-practice perspective, a U.S.-South Korea joint approach can both mitigate AI weaponization risks and help reduce investment uncertainty for companies in allied countries. To this end, a new governance direction centered on compute quotas, access permissions, Zero Trust, and real-time access control can serve as a foundation for establishing future joint working groups and step-by-step policy alignment.

As a practical first step, the two countries should establish a joint working group to develop common standards for access logging, compute-usage monitoring (see Appendix B), and the

identification of high-risk AI functions in cross-border data center environments. A second priority should be to pilot a compute-quota-based governance model in a limited set of high-risk domains—such as military-relevant model training, critical-infrastructure AI systems, or dual-use intelligence applications—before considering broader deployment across other allied countries. These pilots would allow regulators and industry to test different logging architectures, threshold definitions, and enforcement mechanisms in a controlled setting, generating the empirical and institutional experience needed to scale up joint governance of AI infrastructure over time.

On this basis, there are several policy questions that policymakers and practitioners should consider in relation to U.S.-South Korea joint export control. In which domains and for which targets should the shift to compute-centric control be examined first? What criteria (roles, purposes, nationality, risk grades) are needed to treat access-permission structures, rather than physical borders, as the primary units of control? Under what conditions are access control using Zero Trust and real-time access enforcement implementable, and where do technological and legal limitations arise? If compute quotas are introduced as a policy variable, what thresholds, time periods, and project units are appropriate, and how should these be differentiated between allies and countries of concern? What could a U.S.-South Korea joint working group pilot first (for example, log standards, compute monitoring, criteria for identifying high-risk APIs)?

In short, as AI data centers advance, U.S.-South Korea export control laws should also be updated. This is not simply a matter of tightening existing rules, but of redesigning alliance-based national security governance by introducing access and compute as new units of control.

Appendix A. Explanatory Note on the Security Implications of GPU Counts and Compute Quotas

A.1. Why "The Same One Hundred GPUs" Can Mean Very Different Things

In the main text, this paper points out that the security implications differ between a scenario in which one hundred identical GPUs are distributed across multiple allied research institutions and one in which a single organization monopolizes them over a long period. The difference arises not from the physical number of devices, but from how compute is distributed or concentrated.

When 100 GPUs are shared among public research institutes or universities across several allied countries, the compute capacity available to each institution is limited, and research tends to be more diversified. In this case, total compute is distributed across many actors, so the likelihood that any single organization's strategic AI capabilities will leap forward rapidly is relatively low. By contrast, when the same 100 GPUs are concentrated for an extended period in a particular military agency, intelligence service, or state-owned research institute and repeatedly used to train frontier-level models, compute accumulates with a single actor, enabling the rapid development of large-scale models that can be directly applied to military missions such as target identification, situational awareness, and cyber operations.

A.2. "How Many Units" Versus "Who Uses How Much, and for How Long"

Traditional export control asks: "How many GPUs of what performance were exported to which country?" However, AI capability is better explained by the total compute used for a training run and the cumulative compute quota an organization uses over time. The key question thus shifts from "How many GPUs are there?" to "Which actor is using how much compute, for how long?"

A.3. Implications for the Design of U.S.-South Korea Export Control Regimes

This distinction suggests two main directions for the design of U.S.-South Korea export control regimes. First, regulations that focus solely on device counts and performance are insufficient to fully capture actual patterns of compute concentration in cloud-based and multi-tenant environments. Second, if compute quotas and their distribution are introduced as policy variables, it becomes possible to design more fine-grained ceilings on the strategic AI capabilities permitted to allies, friendly states, and countries of concern. Ultimately, even for the same 100 GPUs, security implications and policy responses must vary depending on whose hands they are in, what usage patterns they follow, and how long they are concentrated.

Appendix B. International Regulatory Trends on Logging, Compute Quotas, and Real-Time Control

B.1. Institutionalization of Logging and Record Management: Trends in Europe and the United States

In EU AI regulatory discussions, the EU AI Act (Regulation (EU) 2024/1689) requires providers of high-risk AI systems to automatically generate and retain logs of inputs, outputs, and system behavior, with the aim of ensuring traceability and explainability. The recordkeeping provisions of the U.S. EAR (15 CFR Part 762) impose broad retention duties that extend to electronic records and system-generated logs. For companies that handle controlled technical data in cloud or remote-access environments, the systematic management of access logs—showing who accessed which data, under what conditions, and when—is recognized as a key compliance element.

B.2. Limitations of the U.S. System and Directions for Improvement

Current deemed export rules and the EAR's concept of "release" largely preserve a traditional structure focused on the "transfer of knowledge" and the "disclosure of technology."[28] While this framework works reasonably well for regulating situations in which foreign persons gain access to technology through documents and source code, it is increasingly insufficient to capture AI-related risks that manifest in cloud environments primarily through execution and compute—especially high-risk API-based functional access and large-scale concentration of compute.

Recently, the Trump administration and major research institutions have been discussing ways to impose notification and reporting obligations on training runs for large models that use compute above certain thresholds, and to use such reports as a basis for monitoring AI capabilities. These discussions could be integrated with export licensing conditions so that compute quotas exceeding specific thresholds are subject to more stringent review, reporting, and restrictions. These directions should not be seen as a purely U.S. domestic matter but as issues for joint discussion within a U.S.–South Korea governance framework.

B.3. EAR Provisions on the Transfer ("Release") of Technical Data

Under 15 CFR § 734.15, technology is "released" to a foreign person when that person is permitted to inspect or receive it, including through access information such as passwords or decryption keys. Under 15 CFR § 734.19, transferring such access information may itself constitute a licensable activity. These provisions indicate that existing law already recognizes certain forms of access-enabled release but does not yet provide a sufficiently clear framework for AI-specific compute access scenarios, such as high-risk API-based functional use or large-scale compute concentration.

B.4. ITAR Provisions on Technical Data and Release

Under 22 CFR § 120.33, ITAR broadly defines technical data and treats the provision of access credentials to encrypted defense technical data as a potential release. This framework presupposes the importance of access logs and underscores that providing the means of access—not merely the data itself—can trigger export control obligations. The policy implication is that AI-specific access controls must be designed with this interpretive risk in mind.

B.5. Provisions on Technology Transfer Under South Korean Export Control Law and Their Implications

South Korea's export control regime—anchored in the Foreign Trade Act and the Public Notice on the Export and Import of Strategic Items—defines technology provision primarily in terms of active acts such as transfer of materials, training, and technical guidance. This framework has not yet explicitly captured access-based usage modes such as cloud access, remote API calls, or cumulative compute utilization as distinct regulatory units. Refinement is needed in three areas: 1) clear criteria for when cloud- and API-based usage constitutes strategic technology provision; 2) guidelines for managing access logs and compute-usage records for AI infrastructure; and 3) introduction of compute scale and access patterns as explicit variables in licensing and monitoring frameworks.

Endnotes

[1] Alexandra Alper and Stephen Nellis, "US Mulls New Rules for AI Chip Exports, Including Requiring US Investments by Foreign Firms," Reuters, March 5, 2026, https://www.reuters.com/world/us-mulls-new-rules-ai-chip-exports-including-requiring-investments-by-foreign-2026-03-05; Mackenzie Hawkins, "US Considers Requiring Permits for Nvidia, AMD Global AI Chip Sales," *Bloomberg*, March 5, 2026, https://www.bloomberg.com/news/articles/2026-03-05/us-drafts-rules-for-sweeping-power-over-nvidia-s-global-sales.

[2] South Korean Ministry of Science and ICT, "The AI Basic Act Comes into Force to Lay the Foundation for Korea to Become an AI G3," January 22, 2026, https://www.msit.go.kr/eng/bbs/view.do%3Bjsessionid%3DZT0iXB7mAiF9kdAY5Ak7c74gZdsb4OTVG2h47Huj.AP_msit_1?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1214&sCode=eng.

[3] "S. Korea Begins GPU Rollout to Boost AI Research, Industry," United Press International, March 3, 2026, https://www.upi.com/Top_News/World-News/2026/03/03/gpu-distribution-industry-academia-research-institutions/9711772587666.

[4] Xiaoyang Zhang, "From COCOM to Wassenaar: Is It Still Our Way Ahead?" *Drexel Law Review* 15, no. 1 (2023): 47–119, https://drexel.edu/~/media/Files/law/law%20review/v15-1/Zhang%2047.ashx; "The Wassenaar Arrangement at a Glance," Arms Control Association, last updated February 2022, https://www.armscontrol.org/fact-sheets/wassenaar; *Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies Volume I: Founding Documents* (Wassenaar Arrangement Secretariat, 2019), https://www.wassenaar.org/app/uploads/2019/12/WA-DOC-19-Public-Docs-Vol-I-Founding-Documents.pdf.

[5] "Removing Barriers to American Leadership in Artificial Intelligence," January 31, 2025, https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence; The White House, *Winning the Race: America's AI Action Plan* (2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.

[6] Ministry of Science and ICT, "National AI Strategy Policy Directions," September 26, 2024, https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1040&pageIndex&sCode=eng&searchOpt=ALL&searchTxt; Government of the Republic of Korea, State Affairs Planning Committee, "National Agenda: Five-Year Plan of State Administration," June 2025, https://www.korea.kr/govVision/; "NVIDIA, South Korea Government and Industrial Giants Build AI Infrastructure and Ecosystem to Fuel Korea Innovation, Industries and Jobs," NVIDIA, October 31, 2025, https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-South-Korea-Government-and-Industrial-Giants-Build-AI-Infrastructure-and-Ecosystem-to-Fuel-Korea-Innovation-Industries-and-Jobs/default.aspx.

[7] For examples of cross-border access patterns in multinational data center environments, see major cloud providers' regional infrastructure documentation, e.g., Amazon Web Services, Microsoft Azure, and Google Cloud's regional deployment and access control guides.

[8] European Data Protection Supervisor, "International Transfers," accessed April 28, 2026, https://www.edps.europa.eu/data-protection/data-protection/reference-library/international-transfers_en; "Data Protection Laws in

South Korea," DLA Piper, last updated January 20, 2025, https://www.dlapiperdataprotection.com/index.html?c=KR&t=law; "Understanding Korean PIPA: A Guide for Foreign Businesses," VeraSafe, July 31, 2024, https://verasafe.com/blog/understanding-korean-pipa-a-guide-for-foreign-businesses.

[9] See European Data Protection Supervisor, "International Transfers"; "Data Protection Laws in South Korea," DLA Piper.

[10] U.S. Bureau of Industry and Security, 15 C.F.R. §§ 730–774 (2026), https://www.ecfr.gov/current/title-15/part-730; International Traffic in Arms Regulations, 22 C.F.R. §§ 120–130 (2026), https://www.ecfr.gov/current/title-22/part-120; Foreign Trade Act, art. 19, amended by Act No. 13838 (January 27, 2016), https://elaw.klri.re.kr/eng_service/lawView.do?hseq=37529&lang=ENG; South Korean Ministry of Trade, Industry and Energy, "Public Notice on Export and Import of Strategic Items."

[11] For a definition of "release" of technology, see U.S. Bureau of Industry and Security, 15 C.F.R. § 734.15, https://www.ecfr.gov/current/title-15/section-734.15; Adnan Masood, "Export Controls and Advanced AI Systems in the United States," Medium, February 25, 2026, https://medium.com/@adnanmasood/export-controls-and-advanced-ai-systems-in-the-united-states-ear-itar-ofac-risk-in-models-cloud-35769edcdeaa; Bruce H. Leeds, "Storing Export Controlled Data in the Cloud: What's the Latest?" Braumiller Law Group, accessed April 28, 2026, https://www.braumillerlaw.com/storing-export-controlled-data-in-the-cloud-whats-the-latest.

[12] Masood, "Export Controls and Advanced AI Systems in the United States"; Leeds, "Storing Export Controlled Data in the Cloud: What's the Latest?"; Hanna Dohmen et al., "Controlling Access to Advanced Compute via the Cloud: Options for U.S. Policymakers," Center for Security and Emerging Technology, May 15, 2023, https://cset.georgetown.edu/article/controlling-access-to-advanced-compute-via-the-cloud/.

[13] As of early 2026, neither the U.S. EAR/ITAR nor South Korean strategic technology law explicitly defines API-based access to computation as a distinct regulatory category. See 15 C.F.R. § 734.15; 15 C.F.R. § 734.19; Dohmen et al., "Controlling Access to Advanced Compute via the Cloud: Options for U.S. Policymakers"; South Korean Ministry of Trade, Industry and Energy, "전략물자수출입고시 [Public Notice on Export and Import of Strategic Items]," MOTIE Notice No. 2025-37 (effective Dec. 31, 2025), https://www.law.go.kr/%ED%96%89%EC%A0%95%EA%B7%9C%EC%B9%99/%EC%A0%84%EB%9E%B5%EB%AC%BC%EC%9E%90%20%EC%88%98%EC%B6%9C%EC%9E%85%EA%B3%A0%EC%8B%9C.

[14] Masood, "Export Controls and Advanced AI Systems in the United States"; Leeds, "Storing Export Controlled Data in the Cloud: What's the Latest?"

[15] U.S. Department of Commerce, "Framework for Artificial Intelligence Diffusion," January 15, 2025, https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion; "Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024," 2024 O.J. (L 2024/1689) 1, https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng.

[16] U.S. Department of Commerce, "Framework for Artificial Intelligence Diffusion"; U.S. Cybersecurity and Infrastructure Security Agency, "Executive Order on Improving the Nation's Cybersecurity," accessed April 28, 2026, https://www.cisa.gov/topics/cybersecurity-best-practices/executive-order-improving-nations-cybersecurity; "Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024."

[17] See C. Randall Pratt, "Cloud Computing and Deemed Exports," U.S. Bureau of Industry and Security, Advisory Opinion, January 11, 2011, https://www.bis.gov/media/1352.

[18] For EAR recordkeeping requirements, see U.S. Bureau of Industry and Security, 15 C.F.R. pt. 762, https://www.ecfr.gov/current/title-15/part-762; U.S. Bureau of Industry and Security, *Export Compliance Guidelines: The Elements of an Effective Export Compliance Program* (2017), https://www.bis.gov/sites/default/files/documents/ECP_0.pdf.

[19] U.S. Bureau of Industry and Security, "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification," October 13, 2022, https://www.federalregister.gov/documents/2022/10/13/2022-21658/implementation-of-additional-export-controls-certain-advanced-computing-and-semiconductor; U.S. Bureau of Industry and Security, "Framework for Artificial Intelligence Diffusion."

[20] Foreign Trade Act, art. 19; Ministry of Trade, Industry and Energy, "전략물자 수출입고시 [Public Notice on Export and Import of Strategic Items]," Notice No. 2023-231 (2023).

[21] "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024," Art. 12 (logging obligations for high-risk AI systems), Art. 19 (retention of logs), 2024 O.J. (L 2024/1689) 1, https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng; U.S. Bureau of Industry and Security, "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification"; U.S. Bureau of Industry and Security, "Framework for Artificial Intelligence Diffusion."

[22] U.S. Bureau of Industry and Security, 15 C.F.R. pt. 762 (EAR recordkeeping requirements); U.S. Bureau of Industry and Security, *Export Compliance Guidelines: The Elements of an Effective Export Compliance Program*.

[23] U.S. Bureau of Industry and Security, 15 C.F.R. § 734.15 (defining "release" of technology), https://www.ecfr.gov/current/title-15/section-734.15; 15 C.F.R. § 734.13(b) (deemed export—release of controlled technology to a foreign person within the United States); Masood, "Export Controls and Advanced AI Systems in the United States"; Leeds, "Storing Export Controlled Data in the Cloud."

[24] Scott Rose et al., *Zero Trust Architecture* (National Institute of Standards and Technology, 2020), https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf; Cybersecurity and Infrastructure Security Agency, "Executive Order on Improving the Nation's Cybersecurity," accessed April 28, 2026, https://www.cisa.gov/topics/cybersecurity-best-practices/executive-order-improving-nations-cybersecurity.

[25] Will Shumate et al., *Export Controls on Artificial Intelligence and Uncrewed Aircraft Systems* (RAND, February 2026), https://www.rand.org/pubs/research_reports/RRA3296-1.html; Mark Bromley and Giovanna Maletta, *The Militarization of Technology: Preventing Diversion and Misuse Through Export Controls* (Stockholm International Peace Research Institute, 2025), https://www.sipri.org/sites/default/files/2025-11/rpp_2025_11_miltech.pdf.

[26] U.S. Department of State, 22 C.F.R. § 121 (2026), https://www.ecfr.gov/current/title-22/part-121 (Category VIII: Aircraft and Related Articles, including UAVs); U.S. Bureau of Industry and Security, 15 C.F.R. § 774 (2026), https://www.ecfr.gov/current/title-15/part-774 (ECCN 9A012: UAVs and related systems); "Missile Technology Control Regime Equipment, Software and Technology Annex," Missile Technology Control Regime, https://www.mtcr.info/en/mtcr-annex.

[27] Shumate et al., *Export Controls on Artificial Intelligence and Uncrewed Aircraft Systems*; Bromley and Maletta, *The Militarization of Technology: Preventing Diversion and Misuse Through Export Controls.*

[28] 15 C.F.R. § 734.15 ("release" of technology or software to a foreign person); 15 C.F.R. § 734.13 (export of technology and software); Export Administration Regulations (EAR), 15 C.F.R. pt. 734 (scope of the EAR, including definitions of "release" and "transfer").