CCIA Views on Proposed Enforcement Decree for South Korea's AI Law

On January 21, South Korea passed the AI Basic Act (the Law).[1] The Law's provisions will enter into force on January 22, 2026, and in September 2025, the Ministry of Science and ICT (MSIT) released draft subordinate regulations (the Enforcement Decree) to implement the Law. As MSIT finalizes the Enforcement Decree, it should address serious concerns with the draft text, including those related to transparency obligations, the process for scoping high-impact AI systems and subjecting them to additional obligations, impact assessment requirements, and broader procedural considerations.

Detailed comments are detailed below.

Definitions

In general, MSIT should more precisely define the respective responsibilities of developers and deployers, as the current use of the term "AI business entity" in the Act[2] and draft decrees is overly broad and inconsistently applied, creating legal uncertainty for different actors in the AI ecosystem. Given the complex and layered nature of AI technologies, the Enforcement Decree should clearly differentiate between the obligations of the developer, who builds and provides the system, and those of the deployer, who integrates it into services offered to end-users.

In addition, the Enforcement Decree should explicitly define a "user" as the end user who ultimately receives the final AI product or service, rather than including entities that integrate AI into their own services. The government's own "Explanation of Key Concepts in the AI Framework Act" highlights the confusion that arises from the current approach, which risks misallocating legal responsibilities.

The Enforcement Decree should also clarify liability in cases where a deployer makes a substantial modification to a general-purpose AI system for a high-impact use beyond the developer's intended purpose. In such cases, the deployer, not the original developer, should be deemed the "high-impact AI developer" and assume corresponding obligations. This approach reflects the fundamental principle that the entity controlling the final application is best positioned to assess and manage its risks. Without such a clarification, original developers could face broad and unpredictable liability for downstream modifications they neither intended nor controlled, creating significant legal uncertainty and discouraging the development and distribution of general-purpose AI systems.

[1] https://likms.assembly.go.kr/bill/billDetail.do?billId=PRC_R2V4H1W1T2K5M1O6E4Q9T0V7Q9S0U0.

[2] Article 2.7 of the Law defines "artificial intelligence business entities" as follows. "2.7. The term "artificial intelligence business entity" means any of the following corporations, organizations, individuals, national agencies, etc. which are engaged in a business related to the artificial intelligence industry: (a) An artificial intelligence developer: A person who develops and provides artificial intelligence; (b) An artificial intelligence deployer: A person who provides artificial intelligence products or artificial intelligence services by using artificial intelligence provided by a business entity under item (a)."

Addendum: Effective Date

The AI Act is set to take effect on January 22, 2026. However, the Enforcement Decree and accompanying 400-page set of guidelines are not expected to be finalized before December of this year. To the extent that AI developers and deployers have to develop procedures, train personnel, and integrate technical capabilities into products to comply with this law, such a short period will not provide a reasonable timeframe. Accordingly, MSIT should announce a moratorium on enforcement of these provisions for at least six months, and preferably for one year.

Article 22: Obligations to ensure transparency of AI

Article 22 of the Enforcement Decree requires AI business entities to give users prior notice before offering any product or service that uses high-impact or generative AI, which can be done by labeling the product, including the information in contracts or terms of use, displaying it on a user’s screen, posting it where the service is provided, or using another approved method. Providers may also label AI-generated outputs in human- or machine-readable formats. Such notices must be clear and easily recognizable, taking into account how users perceive information and their age, physical, or social conditions. The obligation does not apply if it is already obvious that AI is being used, if the system is used solely for internal business purposes, or in other cases defined by the Minister based on the nature and use of the product or technology.

MSIT should revise Article 22 to ensure that the transparency obligations apply to AI deployers, not developers, given that the former determine how, where, and for what purpose AI systems are deployed in consumer-facing contexts. Developers, particularly those located outside the jurisdiction, often have no control or visibility over how their models are integrated, customized, or presented to end users, making it both impractical and inequitable to hold them accountable for disclosure and labeling obligations tied to deployment. In many cases, deployers significantly alter a model’s behavior, such as by fine-tuning weights, integrating it into proprietary applications, or designing the user interface through which outputs are delivered, meaning they alone have the operational knowledge necessary to implement meaningful and context-appropriate disclosures. Limiting these requirements to deployers ensures accountability is placed where it is most effective, on the party with direct control over the user experience and the capacity to comply.

In addition, the guidelines provide little clarity on how entities can fulfill the obligations under Article 31(3) of the original Act, particularly regarding what constitutes a “clear and conspicuous” disclosure for audio outputs that are difficult to distinguish from real ones. The current structure risks imposing duplicative obligations, which would be unnecessarily burdensome without enhancing user understanding. To address this, MSIT should streamline these requirements so that a single, clear method of marking can satisfy both provisions, provided that users are also notified, under Article 31(1) of the Act, that the product or service utilizes generative AI.

Moreover, Article 22 should incorporate mutual recognition clauses that acknowledge adherence to equivalent foreign regulatory regimes as satisfying domestic disclosure and notification requirements. Many AI deployers operate across multiple jurisdictions, and requiring them to independently tailor identical transparency mechanisms for each regulatory system would impose unnecessary costs, create inconsistent user experiences, and risk fragmenting compliance approaches. A mutual recognition mechanism would allow regulators to focus on substantive outcomes and incentivize international regulatory convergence.

MSIT should consider revising Article 22.1 as follows:

> *"(1) Before providing products or services using high-impact AI or generated AI (hereinafter referred to as "products, etc."), an AI ~~business entity~~* ***deployer*** *shall give prior notice under Article 31 (1) of the Act by any of the following means:*
>
> *1. Indicate directly on the product, etc., or state in the contract, manual, terms of use, etc.*
>
> *2. Display on a user's screen, terminal, etc.;*
>
> *3. Posting in an easily recognizable manner at a place where products, etc. are provided (including a place within a range reasonably related to the relevant place);*
>
> *4. Other methods recognized by the Minister of Science and ICT in consideration of the characteristics of the products,* ***including compliance with equivalent disclosure or notification requirements under foreign laws or regulations that provide a comparable level of user protection deemed to satisfy the obligations under this Article.***
>
> ***(2) For the purposes of paragraphs (1) and Articles 31(2) and 31(3) of the Act, a single, clear, and conspicuous method of marking or notification shall be deemed sufficient to satisfy multiple disclosure requirements, provided that users are also informed that the product or service utilizes generative AI as required under Article 31(1).”***

Article 23: Obligations to ensure the safety of AI systems

Article 23 defines the scope of AI systems subject to additional safety requirements, including mandatory submission of risk-mitigation and management plans to MSIT, for those trained using a cumulative computation capacity of at least 10^{26} FLOPs. The methodology for calculating cumulative computation capacity is subject to further guidance by MSIT.

While the requirements for such systems are themselves not onerous, and in line with voluntary practices adopted by industry, the use of compute training as a proxy for risk is not effective and risks disproportionately scoping in foreign, especially U.S., AI developers. Basing regulatory thresholds on computation capacity is a flawed and increasingly outdated proxy for risk. As models become more efficient and specialized, they will achieve advanced capabilities with far less compute, rendering such thresholds obsolete. For instance, China’s DeepSeek-V3, comparable in capability to models that Korea would classify as high-risk, would fall outside the scope of these obligations purely because of its lower compute footprint. Moreover, Article 23 compounds these issues by applying compute thresholds to AI systems rather than to underlying models. This approach not only conflates training, deployment, and fine-tuning compute but also forces duplicative risk assessments for multiple systems built on the same model. Finally, tying heightened obligations to compute thresholds risks creating a discriminatory regime that disproportionately impacts U.S. companies while excluding domestic and Chinese competitors, a result that could conflict with Korea's WTO and KORUS commitment to avoid discriminatory treatment. To avoid these outcomes, the Enforcement Decree should abandon computation thresholds altogether and instead assess risk based on AI systems' capabilities

If, however, the Enforcement Decree ultimately retains compute thresholds as a criterion for designating high-risk systems, MSIT should, at a minimum, adopt certain principles. First, as detailed above, it should ensure that the threshold is applied to AI models rather than systems. Second, it should work closely with industry to ensure its design and application reflect technical realities and international best practices. This includes drawing on approaches that enjoy broad industry consensus and support, thereby improving the legitimacy, predictability, and effectiveness of any compute-based framework.[3] Third, the obligation should apply only to base models that exceed the compute threshold, while fine-tuned models derived from those base models should be explicitly excluded from scope, a crucial safeguard against regulatory overreach that could otherwise chill innovation and the development of downstream applications built on foundational models. In addition, Article 23 of the Enforcement Decree should be revised to explicitly codify the three definitional criteria for determining whether an AI system is subject to safety obligations, relocating these core elements from low-tier guidelines into the legally binding text to enhance clarity and legal certainty. Fourth, MSIT should include a provision requiring periodic review of the compute threshold to account for rapid technological advances and evolving model capabilities.

MSIT should consider revising Article 23 as follows:

> *(1) "Artificial intelligence* ***models*** *~~systems~~, the cumulative computation of which used for learning is at least the standard prescribed by Presidential Decree" in Article 32 (1) of the Act means artificial intelligence* ***base models*** *~~systems~~ that meet the standards publicly notified by the Minister of Science and ICT in consideration of the level of development of artificial intelligence technology, the level of risk, etc., and the calculation of which is at least 26 floating points.* ***Fine-tuned models derived from such base models shall be excluded from the scope of this provision.***
>
> *(2) The public notice under paragraph (1) shall include a specific method of calculating the cumulative amount of computation used for learning****, and shall explicitly set out the three definitional criteria for determining whether a model is subject to the safety obligations.***
>
> ***(3) The Minister of Science and ICT shall periodically review and, if necessary, revise the computation threshold prescribed under paragraph (1), in consultation with industry experts, to reflect advances in artificial intelligence technology and evolving risk profiles.***

[3] For example, see https://www.frontiermodelforum.org/updates/issue-brief-measuring-training-compute/.

Article 24: Procedures for verifying high-impact AI

Article 24 establishes the procedure for AI businesses seeking official verification on whether their systems qualify as "high-impact AI" under the law. Businesses must submit a formal request to MSIT, which will assess the application based on factors such as the sector of use, the severity and frequency of potential risks to human safety and fundamental rights, prior assessments, expert committee opinions, and other relevant data. The Minister must issue a decision within 30 days, extendable by another 30 days for complex cases. If a business disagrees with the determination, it may request a re-verification within 10 days, providing reasons and supporting materials. Upon such a request, the Minister, with expert advice, must conduct a renewed assessment and respond within 30 days.

Article 24's definition and verification criteria for high-impact AI should be narrowed to ensure that only systems intended for use in high-risk contexts fall within the scope of regulation. Without this clarification, even foundational or general-purpose AI systems, which may be deployed in a wide variety of benign or low-risk applications, could inadvertently be classified as high-impact, creating excessive compliance burdens and chilling innovation.

Moreover, the current procedural timeline under Article 24, which allows up to 60 days, and in limited cases up to 90 days, for verification, risks creating legal uncertainty and operational delays for companies seeking to comply in good faith. To address this, the Article should introduce either (i) an expedited confirmation process for systems that are clearly unlikely to be classified as high-impact, with a response required within a shorter timeframe (e.g., 15 days), or (ii) a safe harbor provision shielding companies from liability if they proceed based on a reasonable business judgment during the verification period and are later found to have been operating a high-impact AI system.

MSIT should consider revising Article 24 as follows:

> *(1) Where an AI business entity intends to request verification as to whether AI falls under high-impacted AI pursuant to Article 33 (1) of the Act, it shall submit a request for verification in the attached Form to the Minister of Science and ICT.*
>
> *(2) The Minister of Science and ICT shall determine whether a product falls under high-impact artificial intelligence in consideration of the following:*
>
> *1. Whether artificial intelligence is **provided with the intent to be** used in any of the areas referred to in the items of subparagraph 4 of Article 2 of the Act;*
>
> *2. Impact, seriousness, and frequency of risks that may pose to the safety of human life and body and basic rights, and specificity of each area of use;*
>
> *3. Results of prior review as to whether AI falls under high-impact AI under Article 33 (1) of the Act;*
>
> *4. Results of advice, if consulted by a specialized committee pursuant to Article 33 (2) of the Act;*

5. Other data necessary to verify whether AI falls under high-impact AI, which are prescribed by the Minister of Science and ICT.

(3) The Minister of Science and ICT shall reply within 30 days after receiving a request under paragraph (1). In such cases, the period may be extended by up to 30 days, in consideration of the complexity, importance, etc. of products, etc. ***However, for services that are clearly unlikely to be classified as high-impact AI, the Minister shall establish an expedited confirmation procedure with a response period shorter than 30 days.***

(4) An artificial intelligence business entity in receipt of a reply under paragraph (3) may request re-verification by submitting a document (including an electronic document; hereinafter the same shall apply) stating the purpose and reasons for the request for re-verification, along with necessary data, to the Minister of Science and ICT within ten days from the date of receipt of the reply.

(5) Upon receipt of a request for reconfirmation under paragraph (4), the Minister of Science and ICT shall reconfirm whether the relevant product falls under high-impact artificial intelligence after obtaining advice from the expert committee under Article 33 (2) of the Act (hereinafter referred to as the "expert committee"), and reply within 30 days from the date of receipt of the request for reconfirmation.

(6) Where a business entity proceeds with deployment or service provision based on a reasonable judgment that its system is not high-impact AI during the period before receiving a verification result, such entity shall not incur liability for non-fulfillment of high-impact AI obligations if it is subsequently determined that the system is high-impact AI.

Article 26: Business Obligations for High-Impact AI

Article 26 requires AI businesses to publicly disclose key elements of their risk management practices related to high-impact AI systems on their websites, excluding any information considered trade secrets under relevant law. Such disclosures must include core details of risk management policies and structures, standards and explanatory measures, user protection steps, and contact information for the individual responsible for overseeing the AI system. If a deployer uses an AI system whose developer has already fulfilled these obligations and does not significantly alter the system's functions, the AI business entity is considered compliant under Article 26. Deployers may also request necessary information from developers, who are expected to cooperate. Additionally, AI businesses must implement and document these measures for at least five years.

Article 26 should be revised to clarify that the primary responsibility for implementing and disclosing risk management obligations lies with AI deployers rather than developers. Because deployers are the parties that determine how and where AI systems are used, and thus are most familiar with the context, risk profile, and potential impacts of deployment, they are better positioned to conduct appropriate risk management and fulfill obligations under Article 26(1). Assigning these responsibilities to deployers would also ensure that compliance is tied to real-world use cases rather than to the development stage, where many relevant risks may not yet be apparent.

The current requirement to publicly post internal risk management measures should be removed or substantially narrowed. Such a mandate exceeds the scope of the parent law, which does not require publication of internal risk policies or organizational structures. Disclosure obligations should be limited to what is necessary to ensure transparency and accountability without forcing companies to reveal sensitive internal processes that could compromise security or competitive advantage. Regulators could retain the authority to review risk management documentation confidentially as part of compliance checks, but there should be no blanket requirement for public posting.

The Article's current protections, which only exempt "trade secrets" from public disclosure, are too narrow. Companies routinely handle other categories of highly sensitive information, including personal data, commercially sensitive operational information, and security-related details, that should not be exposed. The Article should be amended to explicitly extend confidentiality protections to these categories, ensuring that sensitive data submitted to regulators or included in compliance documentation is safeguarded from disclosure.

MSIT should consider revising Article 26 as follows:

> *(1) An AI ~~business entity~~* ***deployer*** *shall ~~post on its website, etc.~~* ***publicly disclose, where appropriate****, any of the following measures taken under Article 34 (1) of the Act: Provided, That* ***information disclosure shall be limited to information necessary to ensure transparency and accountability****,* ***and*** *matters falling under trade secrets under subparagraph 2 of Article 2 of the Unfair Competition Prevention and Trade Secret Protection Act****, as well as other categories of sensitive information, including personal data, commercially sensitive operational information, and security-related details*** *may be excluded:*
>
> *1.* ***Summary information on*** *key details of the risk management plan under Article 34 (1) 1 of the Act, such as* ***high-level descriptions of*** *risk management policy and organizational system;*
>
> *2. Key details of the standards and explanatory measures under Article 34 (1) 2 of the Act;*
>
> *3. Measures to Protect Users 4. Name and contact information of the person who manages and supervises the relevant high-impact artificial intelligence;*
>
> *(2) Where AI Deployer who has been provided with an artificial intelligence system that has fully or partially performed the measures set forth in Article 34 (1) 1 through 3 of the Act does not cause any significant change in the functions of the artificial intelligence system, it shall be deemed to have performed the measures set forth in Article 34 (1) of the Act.*

(3) The AI Deployer may request the AI Developer to provide necessary data, and the AI Developer shall endeavor to cooperate with such request.

(4) An artificial intelligence ~~business entity~~ ***deployer*** *shall implement the measures set forth in the subparagraphs of Article 34 (1) of the Act and retain the basis thereof in writing for five years.*

(5) "Where the measures prescribed in attached Table 1 have been implemented, as prescribed by Presidential Decree" in Article 34 (3) of the Act means where the measures prescribed in attached Table 1 have been implemented in accordance with the relevant Act.

Article 27: High-Impact AI Impact Assessment

Article 27 outlines the requirements for conducting an AI impact assessment under Article 35(3) of the Act for high-impact AI systems. The assessment must identify the individuals or groups whose fundamental rights could be affected, specify which rights may be impacted, and evaluate the social and economic consequences of those impacts. It should also examine how the AI is used, detail the quantitative or qualitative metrics and calculation methods used in the assessment, and include strategies for risk prevention, loss mitigation, and remediation. If the assessment reveals a need for improvements, it must present the results and an implementation plan. AI companies may carry out the assessment themselves or outsource it to a third party, and MSIT may issue more detailed guidance on assessment methods and procedures.

The current draft of Article 27 does not delineate whether developers or deployers bear primary responsibility for conducting impact assessments. Because deployers determine how and in what context high-impact AI systems are used, they are best positioned to assess potential effects on fundamental rights, societal impacts, and usage behaviors. Article 27 should therefore clearly place primary responsibility on deployers while still allowing them to rely on information provided by developers where necessary.

While deployers should lead the impact assessment process, they often require technical details about the system that only developers can provide. The Article should explicitly state that developers must make reasonable efforts to cooperate with deployers by supplying relevant documentation, data, and technical explanations needed to complete the assessment. This approach balances accountability with practical implementation.

To reduce duplicative compliance burdens and promote regulatory interoperability, the Article should allow deployers to rely on impact assessments already conducted under internationally recognized standards or frameworks (e.g., OECD, ISO/IEC, or other trusted jurisdictional processes), provided that those assessments are substantially equivalent. Such recognition would align with global best practices and support cross-border innovation while still ensuring robust protection of fundamental rights.

MSIT should consider revising Article 27 as follows:

(1) ***An AI deployer shall be primarily responsible for conducting the*** *AI impact assessment under Article 35 (3) of the Act (hereinafter referred to as "impact assessment")****, which*** *shall include the following:*

1. Identification of objects whose fundamental rights are likely to be affected by products or services using the relevant high-impact AI (referring to identification of individuals or groups having certain characteristics);

2. Identification of the types of fundamental rights that may be affected by the high impact AI

3. Details and scope of the social and economic impact on the fundamental rights of a person that may arise from the high-impact AI;

4. Usage behavior of the relevant high-impact AI

5. Quantitative or qualitative evaluation indicators used in the impact assessment and the method of calculating the results;

6. Prevention of risks caused by the high impact AI, recovery of losses, etc.

7. Matters concerning the results of the impact assessment and the implementation plan, if any improvement is required;

(2) An artificial intelligence ~~business entity~~ ***deployer*** *may conduct impact assessment directly or by outsourcing it to a third party****, and may rely on an assessment previously conducted under internationally recognized standards or frameworks, provided such assessment is substantially equivalent to the requirements of this Article.***

(3) The AI developer shall make reasonable efforts to provide relevant technical information, documentation, and cooperation necessary for the deployer to complete the impact assessment.

*(~~3~~**4**) Except as otherwise expressly provided for in the Act and this Decree, the Minister of Science and ICT may establish and disseminate specific details and methods of impact assessment.*

Conclusion

As MSIT finalizes the Enforcement Decree for the AI Basic Act, it is critical that the Decree tailor regulatory obligations to the entities best positioned to fulfill them. Across key provisions, including transparency requirements (Article 22), the verification process for high-impact AI (Article 24), risk management and disclosure obligations (Article 26), and impact assessments (Article 27), primary responsibility should rest with deployers, who determine how and in what contexts AI systems are used, while developers should have a defined duty to cooperate by supplying necessary technical information. The framework should also incorporate mutual recognition mechanisms and acceptance of internationally recognized assessments to reduce duplicative compliance burdens and support cross-border regulatory interoperability. In addition, the scope of “high-impact AI” should be narrowly and clearly defined to target genuinely high-risk contexts, public disclosure requirements should be limited to information strictly necessary for accountability, and protections for sensitive information should be broadened beyond trade secrets to include personal, operational, and security-related data.

To ensure successful implementation, MSIT should also provide longer preparation timelines for businesses to adapt to new requirements and issue clear, detailed regulatory guidance early in the process. During this period, MSIT should commit to a moratorium on enforcing these requirements.

South Korea’s Enforcement Decree for its Basic AI Act has the opportunity, if properly developed, to enhance the country’s position as a leader in global AI governance. Given Korea’s strong domestic AI ecosystem and expanding international market presence, overly restrictive regulations could deter foreign investment, impede R&D, and undermine the competitiveness of Korean companies and startups. A measured, internationally interoperable approach will safeguard fundamental rights while fostering innovation, investment, and long-term economic growth.