

South Korea's AI Basic Act: A Practical Guide for Business Operators

Compliance, Obligations, and Opportunities under Korea's AI Regulation

Kim Kyung-jin, Attorney at Law

AI Library - kimkj.com - 2026

Table of Contents

1. Preface

[Part 1] Why Business Operators Must Know This Act Right Now

- 2. Chapter 1. Legislative History and the 2026 Amendment: What Changes for Business Operators, and From When
- 3. Chapter 2. Am I Subject to This Law?: Definition and Scope of Application of AI Business Operators

[Part 2] Self-Diagnosis: What Regulations Apply to Our Service?

- 4. Chapter 3. Step 1 Determination: Are We a Generative AI Business Operator?
- 5. Chapter 4. Two-Step Determination: Are We a High-Impact AI Business Operator?
- 6. Chapter 5. Three-Stage Assessment: Are We High-Performance, Cutting-Edge AI Business Operators?
- 7. Chapter 6. Step 4 Assessment: Are We an Overseas Business Operator Required to Designate a Domestic Agent?

[Part 3] Basic Obligations for All Business Operators: Securing Transparency

- 8. Chapter 7. Prior Notification Obligation: When, What, and How to Notify
- 9. Chapter 8. Output Labeling Obligations: Watermarking, Labeling, and Special Provisions for Deepfakes
- 10. Chapter 9. Transparency Guidelines in Practice: Closed Networks, Open Networks, and Grounds for Exception

[Part 4] Business Operators Subject to Enhanced Obligations: High-Performance and High-Impact

- 11. Chapter 10. Obligations of High-Performance AI Business Operators to Ensure Safety
- 12. Chapter 11. Five Duties of High-Impact AI Business Operators
- 13. Chapter 12. High-Impact AI Impact Assessment -- De Facto Essential, Not Legally Compulsory
- 14. Chapter 13. Procedure for Requesting Confirmation of High-Impact Status

[Part 5] Adjacent Law Risks Faced by Business Operators

- 15. Chapter 14. Training Data and Copyright: Fair Use Provisions, First Domestic Lawsuits, and Warnings from Overseas Precedents
- 16. Chapter 15. Intersection with the Personal Information Protection Act
- 17. Chapter 16. Overlap with Industry-Specific Adjacent Regulations

[Part 6] Support and Opportunities Available to Business Operators

- 18. Chapter 17. Training Data, Infrastructure, and Aggregation Complexes: Laws Granting Government Support
- 19. Chapter 18. Priority Consideration in Public Procurement and Civil Servant Immunity Clause -- The Key to Opening the B2G Market

20. Chapter 19. Accessibility Support for Vulnerable Groups and Cost Subsidies: A New Market Opened by the 2026 Amendment

21. Chapter 20. Startups, Specialized Personnel, Venture Investment Fund of Funds, and AI Research Institutes -- How to Secure Capital and Talent

[Part 7] In Case of Violation: Sanctions and Risk Management

22. Chapter 21. Initiation Requirements for Factual Investigations and Responding to Document Submission

23. Chapter 22. Administrative Fines, Corrective Orders, Penalties Framework, and the Reality of Grace Periods

24. Chapter 23. Document Management in Preparation for Supervision — What and How Much Should Be Retained?

[Part 8] Overseas Expansion and How to Read Foreign Competitors

25. Chapter 24. South Korean Law Side-by-Side with the EU AI Act

26. Chapter 25. Practical Implications of the 2026 EU Omnibus Revision for Korean Business Operators

27. Chapter 26. US Colorado and Utah State Laws and NIST AI RMF: A Benchmark for Global Business Operators

[Part 9] Practical Guidance for Industry-Specific Business Operators

28. Chapter 27. Financial AI Business Operators

29. Chapter 28. Healthcare and Bio AI Business Operators

30. Chapter 29. Platforms and Generative AI Business Operators

31. Chapter 30. Public Sector Supply Business Operators

[Appendix]

32. Appendix

Preface

On January 22, 2026, the AI Basic Act came into effect in the Republic of Korea. On July 21 of the same year, its amended act also enters into force. Any business operator developing or using artificial intelligence must understand what obligations this Act imposes and what opportunities it creates for them.

This book was written by an attorney for business operators. It is a practical guide, not a legal textbook. It answers questions such as "Is our service high-impact artificial intelligence?", "How should prior notification be conducted?", and "What additional preparations are needed when expanding into the EU?".

There are many issues for which answers are difficult to obtain simply by reading the legal articles. One must examine the Enforcement Decree, public notices, guidelines, and pending follow-up legislation together to see the whole picture. Comparative references to the EU AI Act and Chinese regulations have also been included. This is because Korean companies operating in the global market cannot afford to know only Korean law.

I attempted to write as accessibly as possible. Where legal terms could not be avoided, explanations were placed right next to them. Article numbers have been included without exception, so readers who wish to verify the original text can easily find the corresponding provisions on the National Law Information Center.

This law is in its early stages of enforcement. A grace period of over one year is being operated, and the Enforcement Decree and public notices are being refined sequentially. Some of the interpretations and outlooks presented in this book may change depending on subsequent legislation or official interpretations. While acknowledging these limitations, clear judgments have been set out regarding what business operators must prepare right now.

I hope this book provides practical assistance to artificial intelligence business operators in reducing legal risks and discovering government support and new market opportunities.

July 2026

Kim Kyung-jin

Part 1 Why Business Operators Must Know This Act Right Now

Chapter 1. Legislative History and the 2026 Amendment: What Changes for Business Operators, and From When

On the afternoon of December 26, 2024, electronic voting results appeared on the screen in the main chamber of the National Assembly. Out of 264 members present: 260 in favor, 1 opposed, 3 abstentions. The Framework Act on Artificial Intelligence Development and Trust Building (popularly known as the AI Basic Act) passed the National Assembly. It was the day the world's second law comprehensively regulating artificial intelligence—following the European Union's AI Act—was born. It took four and a half years to reach this vote. Navigating shipwreck in the 21st National Assembly, 19 separate bills were consolidated into one in the 22nd National Assembly, and ruling and opposition party members clashed over single terms in the deliberation room of the Science, ICT, Future Planning, Broadcasting and Communications Committee (hereinafter referred to as the Science and ICT Committee). Even after the bill passed, another amendment was made just two days before it entered into force.

1. A Four-and-a-Half-Year Journey: From Introduction in 2020 to Passage in 2024

In July 2020, the first basic bill on artificial intelligence was introduced in the 21st National Assembly. In July 2021, Representative Pill-mo Jeong and 22 others introduced the "Bill on the Promotion of Artificial Intelligence and Establishment of Trust," followed by additional bills with similar intent introduced by various lawmakers. The issue was the direction. The side advocating for industry promotion and the side advocating for prior safety regulations failed to find common ground. Differences of opinion between the ruling and opposition parties regarding the level of regulation remained unbridged for over three years, and with the expiration of the 21st National Assembly's term in May 2024, all related bills automatically lapsed.

The 22nd National Assembly went into action starting the day after its opening. On May 31, 2024, Representative Cheol-soo Ahn introduced the "Bill on the Promotion of the Artificial Intelligence Industry and Securing of Trust" (Bill No. 2200053). 5 bills were submitted by the ruling People Power Party (including Representatives Jeom-sik Jeong, Seong-won Kim, June-young Bae, and Hee-yong Jeong), 12 bills by opposition parties such as the Democratic Party of Korea (including Representatives In-chul Cho, Hyeong-bae Min, Chil-seung Kwon, Hae-min Lee, and Min-hee Choi),

and 2 bills jointly by ruling and opposition lawmakers. By November 21, 2024, a total of 19 bills had been submitted. The Information, Communication, and Broadcasting Bill Review Subcommittee under the Science and ICT Committee consolidated these 19 bills into a single Committee Chairman's alternative bill, which passed the full Science and ICT Committee meeting on November 26 and was passed with amendments by the Legislation and Judiciary Committee on December 17. On December 26, it was brought to the plenary session and passed with an overwhelming majority.

2. Three Decisions in the Science and ICT Committee Deliberation Room

Three issues were fiercely debated during the process of combining the 19 bills into one.

Korea focuses on impact. The European Union focuses on risk. The EU AI Act uses the term 'high-risk.' Although 'high-risk' appeared in the early stages of the Korean bill as well, concerns were raised that it placed a negative stigma on the technology itself. Representative Hae-min Lee stated, "The reason the global community is moving toward 'impact' is that impact can encompass both risk and scope." The final draft adopted 'high-impact AI.' This means the focus is placed on social repercussions rather than the technology itself.

Whether to prevent beforehand or address after the fact. Early versions of the bill included provisions making prior impact assessment mandatory. Businesses protested, arguing it would erode global competitiveness. A compromise was reached. The impact assessment for high-impact AI was designated as a 'voluntary effort obligation,' but structured so that business operators who complete the impact assessment receive priority consideration in public procurement. Although the law does not compel it, it is a 'soft regulation' under which high-impact AI must effectively undergo assessment to be supplied to the public sector.

After much debate, the level of administrative fines was set low. The maximum administrative fine under the EU AI Act is 35 million euros or 7% of global annual turnover. The maximum administrative fine under Korea's AI Basic Act is 30 million KRW. This reflects the rationale that, because the domestic AI industry is in its early stages of growth, excessive penalties would stifle innovation. Questions over whether the amount of 30 million KRW can provide large enterprises with an incentive to comply continue to be raised even after the law took effect.

3. 2025: Enforcement Decree and Guidelines Are Formulated

On January 21, 2025, the AI Basic Act was promulgated as Act No. 20676. Its enforcement date was set one year later, on January 22, 2026. The intervening year was spent formulating the Enforcement Decree.

The Ministry of Science and ICT established a 'Subordinate Legislation Task Force' comprising about 80 private-sector experts, which held over 70 meetings. On September 8, 2025, the draft Enforcement Decree was reported at the 1st plenary meeting of the National AI Strategy Committee. On the 17th of the same month, drafts of 2 public notices and 5 guidelines were also disclosed. The advance notice of legislation for the Enforcement Decree ran for 40 days from November 12 to December 22, collecting opinions from industry, academia, and civil society.

On January 21, 2026, the Enforcement Decree (Presidential Decree No. 36053) was promulgated. It consists of 33 articles in the main text and 2 addenda articles.

The criteria for determining high-impact AI were established as a two-stage process. In Stage 1, the domain is verified. Fields such as energy, healthcare, recruitment, loan screening, transportation, public services, and student evaluation are covered. In Stage 2, it is

evaluated whether there is a concern of 'having a significant impact on or posing a risk to human life, bodily safety, and fundamental rights.' Guidelines recommend documenting the grounds for determination, and if a request to confirm applicability is submitted to the Minister of Science and ICT, a reply is received within a basic 30-day period, plus an additional 30 days if extended.

The threshold for high-performance AI (large-scale AI) was also established. Systems subject to this threshold are those with cumulative training compute of 10 to the 26th power FLOPs (10^{26} FLOPs) or higher that apply state-of-the-art technology and pose concerns of widespread impact on life, bodily safety, and fundamental rights. Compared to the EU AI Act setting 10 to the 25th power FLOPs (10^{25} FLOPs) as its threshold, this is a threshold one order of magnitude higher. Most domestic business operators do not fall into this bracket.

On the following day, January 22, the Act and the Enforcement Decree entered into force simultaneously. The Ministry of Science and ICT distributed five major guidelines on the same day: Guidelines on Ensuring Transparency, Guidelines on Ensuring Safety, Guidelines on Determining High-Impact AI, Guidelines on Business Operator Duties,

and Guidelines on AI Impact Assessment. Although they lack legal binding force, they serve as practical guidance showing business operators how to fulfill their statutory obligations. Two public notices (Notice on Ensuring the Safety of State-of-the-Art Artificial Intelligence and Notice on the Duties of High-Impact Artificial Intelligence Business Operators) also came into effect together.

On the same day, the AI Basic Act Support Desk opened. Operated by the Korea Software Industry Association (KOSA), a total of 552 consultations were received as of March 31, 2026, the 10th week of implementation. Anonymous consulting is available, and in-depth legal reviews take up to 14 days.

4. Amended Before Implementation: The Amended Act of January 20, 2026

The original AI Basic Act placed its center of gravity on safety and regulation. Promotion provisions such as industry fostering, public sector adoption, and support for vulnerable groups were insufficient. The 22nd National Assembly promulgated the amended Act two days before the original Act took effect.

During 2025, Representatives Min-hee Choi, Jeong-heon Lee, Cheol-min Jang, Bo-yoon Choi, and others sponsored 9 amendment bills. The Information, Communication, and Broadcasting Bill Review

Subcommittee examined them on November 21 and created a Committee alternative bill by consensus between the ruling and opposition parties. It passed the full Science and ICT Committee meeting on November 24 and passed the plenary session on December 30. It was promulgated as Act No. 21311 on January 20, 2026.

The amendments consist of eight items.

The National Artificial Intelligence Committee was renamed the National Artificial Intelligence Strategy Committee (Article 7). It was not merely a name change. Its authority over policy coordination and performance management was expanded. The President serves as Chairperson, the total number of committee members is within 60, and one of the three Vice-Chairpersons serves full-time. With 18 ex-officio members (heads of 16 central administrative agencies, 3rd Deputy Director of the National Security Office, and Senior Secretary to the President for AI), up to 41 appointed private-sector members, plus up to 52 subcommittee members across 8 subcommittees, it forms an 85-member structure in total. Chief AI Officers (CAIOs) of all 48 central administrative agencies are designated at the vice-ministerial level and hold regular monthly consultative

meetings.

Provisions promoting the adoption of AI in the public sector were newly added (Article 16 (3) and (4)). State organs must give priority consideration to AI when purchasing products or services or awarding service contracts. Immunity provisions for public officials were also included. In the absence of intention or gross negligence, public officials bear no liability for damages arising from AI adoption. From a corporate perspective, this structure opens up new demand in the business-to-government (B2G) market.

Legal grounds for establishing artificial intelligence research institutes were established (Article 22-2). Universities or businesses can establish them upon obtaining permission from the Minister of Science and ICT.

Legal grounds for supporting vulnerable groups were established. The amended Act introduced the new concept of vulnerable groups experiencing difficulty using artificial intelligence, and specified legal grounds for financial support regarding access to high-cost AI services (Article 17-2). The amended Enforcement Decree promulgated on July 20, 2026 defined its scope to include persons with disabilities, senior citizens aged 65 or older, recipients of

National Basic Living Security benefits and the near-poor class, job seekers, farmers and fishermen, members of multicultural families, North Korean defectors, women with career breaks, and eligible recipients of single-parent family support.

In addition, support for startups in the AI field, cultivation of professional talent, legal grounds for providing public data as training data, and support for AI technology education were included in the amended Act.

The effective dates are split into two. The reorganization of the Strategy Committee and support for professional talent apply immediately starting January 22, 2026. Provisions requiring the refinement of subordinate legislation—such as public sector AI adoption, procedures for establishing AI research institutes, and financial support for vulnerable groups—take effect on July 21, 2026.

5. Business Operator Timetable: What Must Be Complied With Now and What Remains Deferred

There are obligations that apply immediately starting January 22, 2026.

First is the prior notification obligation (Article 31 (1)). When providing products or services utilizing high-impact artificial

intelligence or generative artificial intelligence, business operators must inform users of that fact in advance. The generated output labeling obligation (Article 31 (2) and (3)) has also commenced. Outputs created by AI must be labeled according to their respective types. In the case of deepfakes, invisible watermarks alone are insufficient, and visible labeling is essential. Exceptions are granted when it is obvious that the output is AI-generated and for internal operational use.

For high-impact AI business operators, five duties are consolidated and imposed under a single provision, Article 34 (1). They are: establishment and operation of risk management measures (Item 1); establishment and implementation of explanation measures regarding final results, key criteria of judgment, and training data overviews (Item 2); establishment and operation of user protection measures (Item 3); securing a framework for human control and supervision (Item 4); and creation and 5-year retention of documentation evidencing compliance (Item 5). The main contents of each measure must be published on the website pursuant to the Enforcement Decree.

High-performance AI business operators must identify, evaluate, and mitigate risks pursuant to the obligation to ensure safety (Article 32) and submit the results.

Overseas AI business operators without an address or place of business in Korea must designate a domestic agent if they exceed certain criteria such as user count or revenue (Article 36). Failure to comply will result in an administrative fine of up to 30 million KRW.

A grace period applies to the imposition of administrative fines (Article 43). The government announced a policy of deferring enforcement for at least one year, until after January 2027. Administrative guidance and corrective recommendations will take priority in matters such as omissions of transparency obligations or inadequate documentation. However, in severe cases such as deepfake criminal exploitation, fatal accidents, or large-scale human rights violations, immediate fact-finding investigations can be conducted.

Business operators already taking equivalent measures pursuant to existing laws such as the Personal Information Protection Act, the Medical Devices Act, or the Use and Protection of Credit Information Act are recognized as having fulfilled the duties of high-impact AI

business operators (Article 34 (1)) to the extent of overlap. This recognition does not extend to transparency obligations—such as prior notification and generated output labeling—or to safety assurance obligations. It is a mechanism to reduce the burden of double regulation.

6. European Union: Departed First, But Arrived Later

The European Commission officially presented its AI Act proposal on April 21, 2021. This was 9 months after the first bill was introduced in Korea's 21st National Assembly. Europe set the record for introducing the world's first comprehensive AI regulatory legislation. In February of the previous year, 2020, the European Commission had already published the "White Paper on Artificial Intelligence" and established a risk-based approach.

Following the European Commission proposal, a legislative process lasting approximately 3 years and 1 month ensued. In December 2022, the Council of the EU adopted a general approach, and in June 2023, the European Parliament adopted its negotiating position (499 in favor, 28 opposed, 93 abstentions). Trilogue negotiations began in the same month, reaching a provisional agreement after a three-day marathon negotiation from December 6 to 9. Following the final plenary vote in the European Parliament on March 13, 2024 (523 in

favor, 46 opposed, 49 abstentions), unanimous approval by the Council of the EU on May 21, and publication in the Official Journal of the EU on July 12, it entered into force on August 1.

Entry into force differs from application. The EU AI Act has a phased application structure. On February 2, 2025, obligations regarding prohibited AI systems (such as social scoring and subliminal manipulation AI) and AI literacy applied first. The EU possesses something absent in Korea: absolute prohibition provisions. Types such as social scoring, untargeted biometric surveillance, and subliminal manipulation are not permitted at all. On August 2, 2025, obligations regarding general-purpose AI (GPAI) models commenced. The GPAI Code of Practice, drawn up with the participation of around 1,000 stakeholders, received an adequacy decision around the same time. Transparency (documentation, 10-year retention), copyright (respecting robots.txt), and safety and security (systemic risk management) form its three pillars.

High-risk AI obligations were originally scheduled to apply on August 2, 2026. However, on November 19, 2025, the European Commission proposed the 'Digital Omnibus' amendment. The core of this amendment is deadline extension. The application date for

standalone high-risk AI (Annex III) was pushed back by about 16 months, from August 2, 2026 to December 2, 2027. High-risk AI embedded in regulated products (Annex I) was delayed by one additional year, from August 2, 2027 to August 2, 2028. The obligation to operate AI regulatory sandboxes was also postponed to August 2, 2027. An explicit goal to reduce regulatory compliance costs for small and medium-sized enterprises (SMEs) by 35% was specified, and new prohibitions were added regarding non-consensual intimate image generation AI and child sexual abuse material generation AI.

This Omnibus amendment went through a provisional agreement on May 7, 2026, official adoption by the European Parliament on June 16, and final approval by the Council of the EU on June 29.

Korea's AI Basic Act has been applying major business operator obligations—such as prior notification, generated output labeling, and high-impact business operator duties—since January 22, 2026. The high-risk AI obligations of the EU AI Act will only apply in December 2027. A reversal has occurred wherein Europe, having introduced its bill first, lags nearly two years behind Korea in regulating high-risk AI business operators. Korea is the first country to enforce business operator obligations under a comprehensive AI law

on such a broad scale.

The gap in financial penalties is stark. Between the EU's 35 million euros (approx. 55 billion KRW) and Korea's 30 million KRW lies a gap of more than a thousandfold. The EU employs an ex-ante control mechanism that mandates ex-ante conformity assessments. Korea employs a mechanism that induces voluntary market compliance through voluntary impact assessments linked with public procurement. Regarding the high-performance AI compute threshold as well, there is a ten-fold difference between the EU's 10^{25} and Korea's 10^{26} , making the scope of regulated entities fundamentally different.

7. China: Rules First, Comprehensive Law Later

China did not create a comprehensive AI law all at once. Its approach has been to issue sector-specific regulations first, enforce them, and prepare a comprehensive law in the background.

On March 1, 2022, the Provisions on the Management of Algorithmic Recommendations in Internet Information Services entered into force. Jointly promulgated by four government bodies—the Cyberspace Administration of China (CAC), the Ministry of Industry and Information Technology, the Ministry of Public Security, and the

State Administration for Market Regulation—the provisions address transparency and user choice in algorithmic recommendation services. On January 10, 2023, the Provisions on the Management of Deep Synthesis (Deepfakes) entered into force, mandating visible labeling for synthetic content. The Interim Measures for the Management of Generative AI Services, which took effect on August 15 of the same year, stipulated the obligation to adhere to 'core socialist values' and required ex-ante security assessments for services capable of shaping public opinion. Incitement to subvert state power, terrorism, and ethnic hatred content are prohibited. Neither Korea nor the EU has these types of content control provisions.

The ex-ante filing system serves as the core enforcement mechanism. Services with the capacity to shape public opinion or mobilize society must register their algorithms within 10 business days after deployment. As of November 2025, 611 generative AI services have completed filing, and total algorithm filings have exceeded 5,000 cases.

On September 1, 2025, the Measures on Labeling AI-Generated Synthetic Content took effect, imposing a dual labeling obligation

(visible labels along with metadata or digital watermarks). Mandatory national technical standards (GB 45438-2025) were even enacted. On July 15, 2026, the Interim Measures for Anthropomorphic AI Interactive Services will take effect. Earlier in May, Implementation Opinions on the Normative Application of AI Agents was issued as a policy document. Providing virtual intimate relationship services to minors is restricted, and safety evaluations are required for services exceeding a certain scale.

Enforcement is already vigorous. In the first half of 2024 alone, Chongqing saw 11 administrative penalties, the closure of 142 websites, and corrective guidance talks with 101 platforms. In February 2025, the CAC removed generative AI service apps and mini-programs that had not undergone security evaluations. In June 2025, Shanghai's 'Huangpu Sharp Sword' campaign resulted in the deletion of 820,000 illegal content items and the deactivation of approximately 2,700 AI agents.

A comprehensive artificial intelligence law has not yet been enacted. Although the State Council's 2017 'Next Generation Artificial Intelligence Development Plan' declared that the AI legal system would be fully established by 2025, as of July 2026, the official draft of

a comprehensive law has not been publicly released. In March 2024, Chinese legal scholars published a scholar draft (totaling 9 chapters and 96 articles), and Georgetown University's CSET released an English translation. In May 2026, the State Council stated in its annual legislative plan that it would "accelerate comprehensive legislation for the healthy development of artificial intelligence." Anticipated provisions include recording all training data sources, mandatory safety impact assessments for models exceeding 100 billion parameters, human intervention requirements for decisions affecting legal rights, and fines of up to 50 million RMB (approx. 9.5 billion KRW). Approval by the National People's Congress is expected in late 2027, with enforcement expected in 2028.

The regulatory intensity of the three jurisdictions follows a clear order: the EU, China, and Korea. Looking strictly at maximum penalties under current law, the EU stands at 35 million euros, Korea at 30 million KRW, and China is projected to reach up to 50 million RMB (approx. 9.5 billion KRW) once its comprehensive law is enacted. Regulatory approaches also differ: the EU uses a comprehensive regulatory law, China uses a combination of sector-specific regulations and content controls, and Korea uses a hybrid of promotion and regulation laws. China's dual-driven model of 'law

plus standards'—wherein the National Standardization Technical Committee (TC260) drafts soft law that transitions into hard law—has no counterpart in either Korea or the EU.

8. 552 Questions

On July 21, 2026, the remaining provisions of the amended Act entered into force. Priority consideration for public sector AI adoption, procedures for establishing AI research institutes, and financial support for vulnerable groups gained legal effect. Following an advance notice of legislation from May 21 to June 19, 2026, the amended Enforcement Decree was promulgated as Presidential Decree No. 36506 on July 20 and took effect on July 21.

On January 22, 2026, the day the law took effect, Legislative Research Officer Joon-hwa Jeong of the National Assembly Research Service published 'Administrative Tasks for Realizing National AI Strategy.' Refining the concept and requirements of high-impact AI and differentiating obligations by business operator type were identified as remaining tasks. Research Fellow Won-june Jeong of the Korea Legislation Research Institute also pointed out in a September 2025 report that the scope of regulated entities (addressees of the law) remained unclear.

Within 10 weeks of implementation, 552 consultation requests flooded into the AI Basic Act Support Desk. The Act was passed, the Enforcement Decree was issued, and the guidelines were distributed. Yet business operators continue to ask: whether their company's AI service constitutes high-impact AI, in what format prior notification must be provided, and what the criteria are for appointing domestic agents for overseas partners. Between having a law and being able to comply with it, 552 questions remain stacked.

Chapter 2. Am I Subject to This Law?: Definition and Scope of Application of AI Business Operators

January 22, 2026, is the day the AI Basic Act came into effect. The department of radiology at a university hospital in Seoul. A specialist brings up a chest CT image on the monitor. On one side of the screen, an AI reading assistance system marks a suspected pulmonary nodule area with a red circle. This doctor reads dozens of images a day together with AI. Is this hospital an AI business operator as defined in Article 2, Item 7 of the AI Basic Act? Under Korean law, it is not. Because it merely receives the AI reading system to use for medical treatment, this hospital is a user under Item 8. It bears no legal obligations. If the same doctor were to use the same system at a hospital in Munich, that hospital would be a Deployer under the European Union AI Act. It must operate the system in accordance with instructions for use, ensure human oversight, and retain system operation logs. It is the same technology and the same action, yet the legal status is reversed.

1. AI as Viewed by Law

The official title of the AI Basic Act is the "Basic Act on the Development of Artificial Intelligence and Creation of Trust Base, etc." Act No. 20676. It passed the plenary session of the National

Assembly on December 26, 2024, was enacted and promulgated on January 21, 2025, and came into force on January 22, 2026. Two days before its enforcement, on January 20, 2026, provisions including Article 2 were amended; its Enforcement Decree (Presidential Decree No. 36053) was promulgated on January 21, 2026, amended on July 20, 2026, and entered into force the following day.

Article 2 contains twelve definitions. Six of them are essential to business operators.

Item 1 defines "artificial intelligence." "That which implements intellectual capabilities possessed by human beings, such as learning, reasoning, perception, judgment, and language comprehension, through electronic methods." This definition lists five human cognitive functions and encompasses the entirety of technology that electronically emulates them. In academia, this definition is assessed as an anthropomorphic approach that describes artificial intelligence by comparing it to human capabilities.

Article 3(1) of the European Union AI Act takes a different approach. Starting with the term "machine-based system," it describes a system that receives inputs, infers, and generates outputs such as predictions, recommendations, or decisions. Korea views it based on

human capabilities. The European Union views it based on machine functions.

Item 2 separately defines an "artificial intelligence system." "An artificial intelligence-based system that infers outputs such as predictions, recommendations, or decisions that influence real and virtual environments for a given set of objectives, with varying levels of autonomy and adaptiveness." If artificial intelligence in Item 1 is the technology itself, the artificial intelligence system in Item 2 is an operating framework encompassing that technology. It must possess autonomy, possess adaptiveness, and generate outputs through inference.

Korean law has thus adopted a two-tiered structure separating "artificial intelligence" and "artificial intelligence system." The EU AI Act defines it singularly as an "AI system."

This distinction is immediately useful in practice. An RPA or rule-based repetitive automation program used by a logistics company in its warehouse is not an artificial intelligence system. Software that repeats tasks according to predetermined rules possesses neither autonomy, adaptiveness, nor inference. On the other hand, if that logistics company uses a machine learning model

that optimizes delivery routes, that is an artificial intelligence system. It learns from data, adapts to environmental changes, and recommends routes through inference.

2. Development Business Operator, Utilization Business Operator, User

Article 2, Item 7 of the Act divides AI business operators into two categories.

A development business operator under Subitem (a) is "a person who develops and provides artificial intelligence." A utilization business operator under Subitem (b) is "a person who provides artificial intelligence products or artificial intelligence services by utilizing artificial intelligence provided by a business operator under Subitem (a)." A user under Item 8 is "a person who is provided with artificial intelligence products or artificial intelligence services."

The key is the verb "provide." A development business operator provides artificial intelligence itself. A utilization business operator uses that artificial intelligence to provide a product or service. A user merely receives the provision.

OpenAI develops a GPT model and provides it via API. It is a development business operator. A domestic fintech company calls

that API to build a customer consultation chatbot and supplies it to a bank. It is a utilization business operator. The bank receives that chatbot and uses it for customer service. If used solely for internal operations, it is closer to a user. However, if it provides even a finished chatbot to customers under its own name, whether it falls under a utilization business operator pursuant to Article 2, Item 7, Subitem (b) must be examined.

This is the design choice of Korean law. It did not impose obligations on entities that adopt AI and use it for their inherent operations.

Whether profit-seeking or non-profit is irrelevant. Regardless of whether an entity is a corporation, an organization, an individual, or a state organ, any person engaged in economic activities related to the artificial intelligence industry is subject to the Act. If a university laboratory creates an AI model and uses it solely for academic papers, it is not a business operator. If it opens that model as an API and continuously provides it externally, it is a business operator.

An entity can hold both statuses simultaneously. If a financial company develops its own AI credit scoring model and provides it to an affiliated bank, while simultaneously importing external AI to use

for internal loan screening, it is a development business operator with respect to the credit scoring model and a utilization business operator with respect to the loan screening AI.

A development business operator bears primary responsibility. This is the fundamental accountability for the model's design, training, and safety. A utilization business operator bears secondary responsibility. It must provide notice of transparency in the context of its own service and establish risk management measures when applying AI to high-impact domains. If the development business operator has already implemented measures concerning risk management, explanation, and user protection, and the utilization business operator has not significantly altered the original purpose or use of the artificial intelligence system, the utilization business operator is also deemed to have implemented them solely with respect to those three measures. This is a mechanism to prevent double burdens.

Fine-tuning blurs this boundary.

There are companies that take external models, fine-tune them with their own data, and provide the output as a new service. They did not create the original model from scratch. Nor are they using it as is.

There is room for providing services after fine-tuning to fit the statutory phrasing of "a person who develops and provides artificial intelligence." If fine-tuning is an act that substantially alters the model's behavior and the company releases the output under its own name, this company is highly likely to be classified as a development business operator. An official statutory interpretation from the Ministry of Science and ICT has not been issued as of July 2026. Uncertainty remains.

In the first ten days after enforcement, the AI Basic Act Support Desk received 172 inquiries. Transparency obligations accounted for 56.4 percent, confirmation of high-impact artificial intelligence accounted for 17 percent, and definitions accounted for 10.6 percent. This means business operators are asking before this law, "Am I subject to this law?" Cases where the line is unclear—such as whether a cosmetics company providing an advertising service using AI is a business operator or a user, or to what extent a credit scoring company providing AI models to financial institutions is subject to regulation—are accumulating in practice.

3. Generative, High-Impact, High-Performance

On top of the framework of AI business operators, three classifications are superposed.

Article 2, Item 5 is generative artificial intelligence. "An artificial intelligence system that imitates the structure and characteristics of input data to generate text, sound, pictures, video, and various other outputs." ChatGPT, Midjourney, and Suno AI fall under this category. Business operators providing generative artificial intelligence are subject to two additional obligations: marking outputs to indicate they were generated by AI, and providing prior notification to users.

Article 2, Item 4 is high-impact artificial intelligence. "An artificial intelligence system that has a significant impact on or is likely to cause risk to human life, physical safety, and fundamental rights," which is used in domains listed by the Act.

There are eleven domains. Five consist of energy supply, drinking water production processes, healthcare systems, medical devices and digital medical devices, and safety management of nuclear materials and nuclear facilities. Another five consist of analysis and utilization of biometric information such as faces, fingerprints, and irises for criminal investigation or arrest duties; judgments significantly affecting individual rights and obligations such as hiring or loan screening; major operations of transport vehicles, transportation facilities, and transportation systems; qualification

screening and cost collection by the state, local governments, and public institutions; and student evaluation from early childhood education to secondary education. The last one is an open-ended item that can be added by Presidential Decree.

The framework for determining whether an AI is high-impact consists of two steps: Is it used in the relevant domain? And does it have a significant impact on life, physical safety, or fundamental rights? The Korea Legislation Research Institute pointed out that in this determination framework, "impact" and "risk" are used concurrently within the law, and "safety" and "fundamental rights" are linked by "and," leading to confusion in interpretation.

High-performance artificial intelligence does not appear in Article 2 of the Basic Act. Entrusted by Article 32 of the Basic Act, it is prescribed in Article 24 of the Enforcement Decree. There are three criteria, all of which must be met: The cumulative computation used for training must be 10^{26} FLOPs or more; it must be constructed and operated by applying state-of-the-art technology among currently utilized artificial intelligence technologies; and its level of risk must pose a concern of having a widespread and significant impact on human life, physical safety, and fundamental rights. Crossing the

computation threshold alone does not make an AI subject to regulation. Business operators of artificial intelligence systems meeting these criteria are imposed with obligations for identifying, assessing, and mitigating risks throughout the lifecycle, establishing a risk management system to monitor and respond to safety incidents, and submitting the results of such implementation.

The computation threshold is ten times higher than that of the European Union. The computation benchmark used by the EU AI Act for general-purpose AI models with systemic risk is 10^{25} FLOPs. Korea set it at 10^{26} . An intention to narrowly confine the targets of regulation can be discerned.

4. A Law Crossing Borders

This law does not apply only within Korea.

It is Article 4, Paragraph 1: "This Act shall apply even to acts performed outside the territory of the Republic of Korea if they affect the domestic market or users." If an entity provides Korean language services, conducts business targeting Korean consumers, or affects the domestic market in any form, this law follows.

Among AI business operators headquartered overseas, those above a certain scale must designate a domestic agent. Entrusted by Article 36, Article 29, Paragraph 1 of the Enforcement Decree sets four criteria, and fulfilling even one triggers the obligation: Total sales of 1 trillion KRW or more in the preceding year; sales in the artificial intelligence service sector of 10 billion KRW or more in the preceding year; a daily average of 1 million or more domestic users over the immediately preceding 3 months; and business operators that have been subjected to administrative fines for failing to comply with suspension orders or corrective orders. The domestic agent submits results of compliance with safety assurance obligations and requests for confirmation of high-impact artificial intelligence status on behalf of the business operator, and takes charge of support necessary to fulfill high-impact duties.

Article 2 of the European Union AI Act subdivides extraterritorial application into seven types: providers placing AI on the market in the EU, deployers located in the EU, third-country business operators whose outputs are used within the EU, importers, distributors, manufacturers integrating AI into their own products, and affected persons within the EU. China has no provisions explicitly stipulating extraterritorial application. It merely states that it applies when

"providing services to the public within the territory of China," replacing jurisdiction with technical blocking instead of explicit extraterritoriality.

Korea encompasses it in a single sentence. The EU breaks it down by type. China lacks an explicit clause altogether.

5. Exception for Military Purposes

Article 4, Paragraph 2 prescribes an exception.

"This Act shall not apply to artificial intelligence developed and utilized exclusively for national defense or national security purposes, as prescribed by Presidential Decree."

The key is the phrase "exclusively for... purposes." In English, exclusively. Only artificial intelligence created solely for defense and used solely for defense falls under this exception. Dual-use artificial intelligence targeting both military and civilian domains is not an exception. The law applies.

Article 2 of the Enforcement Decree narrows the specific scope. Construction of defense information and communications networks and weapon system operation designated by the Minister of National

Defense; early warning systems, security projects, counter-terrorism activities, and investigations into core technology leaks designated by the Director of the National Intelligence Service; and counter-terrorism and investigation into violations of advanced strategic technologies designated by the Commissioner General of the National Police Agency. Artificial intelligence outside this scope remains subject to the law even if used by the military or intelligence agencies.

Article 2(3) of the EU AI Act employs the same structure: "exclusively for military, defence or national security purposes." Both use an exclusive qualifier and reach the same conclusion that dual-use is regulated. If there is a difference, the EU provision applies in and of itself without separate subordinate legislation, whereas Korea delegated the specific scope to Presidential Decree.

6. The Mirror of the European Union

The Korean AI Basic Act and the European Union AI Act gave different answers to the same question.

They differ right from the classification of business operators. The EU distinguishes between Provider and Deployer. Korea distinguishes between development business operator and utilization business

operator. Although the names appear similar, their scopes differ significantly.

An EU Provider is a person who develops an AI system or a general-purpose AI model, or has one developed, and places it on the market under its own name or trademark. By explicitly specifying "under its own name or trademark," it encompasses even commissioned AI under white-label arrangements. Korea's definition of development business operator has no corresponding provision.

The difference between a Deployer and a utilization business operator is even greater. An EU Deployer is any natural or legal person using an AI system under its authority in a professional context. Only personal, non-professional use is excluded. The scope is broad. A Korean utilization business operator is a person who provides products or services using artificial intelligence provided by a development business operator. It presupposes a supply relationship and requires external provision. The scope is narrow.

Let us look again at the department of radiology at a university hospital in Seoul. When a chest CT is taken, the AI reads the image first and leaves a mark on suspected nodules. The doctor writes a reading report with reference to that mark, and the diagnosis is

issued under the doctor's name. The same equipment and the same software are installed at a hospital in Munich, Germany. The machine is the same, but the legal statuses of the two hospitals differ.

In the EU, this hospital is a Deployer. Pursuant to Article 26 of the EU AI Act, it must comply with the manufacturer's instructions for use, assign natural persons having the necessary competence and authority to oversee the AI, and retain system-generated logs for at least six months. In Korea, the same hospital is a user. The AI Basic Act imposes no obligations whatsoever on users.

Why did such a difference arise? The answer lies in the definition clauses. An artificial intelligence utilization business operator is a person who provides "artificial intelligence products or artificial intelligence services" by utilizing AI provided by a development business operator. It does not apply to just any service; what is offered to others must itself be an AI product or AI service. What a hospital offers to a patient is medical treatment. A patient is not provided with an AI service, but with a doctor's diagnosis informed by AI. AI is a tool used and concluded within the hospital. Thus, the hospital is not a provider of AI but a recipient of provision, namely a user. It is the same logic as a bank that imports loan screening AI and

uses it for screening remaining merely a user.

However, there is a boundary line. The criterion for determination is what is handed to the patient. If it is a doctor's diagnosis, the hospital is a user; if it is the output of AI, it is a utilization business operator. If a hospital attaches an AI medical intake chatbot to a patient app so that patients interact directly with the AI, or notifies patients of AI reading results without physician review, what the patient receives becomes the artificial intelligence service itself. At this point, the hospital becomes a utilization business operator with respect to that service, and obligations such as confirming high-impact artificial intelligence status, measures to maintain safety and reliability, and prior notification obligations ensue. One must remember that status attaches not on an institutional basis, but on a service-by-service basis. A hospital that is a user in the radiology department becomes a utilization business operator with respect to that chatbot the moment it launches a patient chatbot.

This structure is a choice of legislation. The EU imposes obligations starting from the stage where a hospital uses AI as a tool, whereas Korea begins obligations only at the stage where a hospital offers AI to a patient. The location of the threshold where regulation

intervenes is different. Korean law relieved the burden on entities adopting AI. The Ministry of Science and ICT itself stated upon issuing guidance guidelines that it gave joint consideration to both the legislative intent of securing AI trust and the easing of corporate burdens.

In practice, determine as follows: If a hospital wishes to maintain user status, specify internal procedures where a doctor makes the final judgment and record it in medical records. If planning to launch an AI service dealing directly with patients, prepare for the obligations of a utilization business operator not on the day of launch, but from the day of preparation. The threshold is low, but the moment it is crossed, the burden is no lighter than that of an EU Deployer.

The EU divides regulated entities into six types: Provider, deployer, importer, distributor, authorized representative, and product manufacturer. These six are collectively called operators. Korea has two: development business operator and utilization business operator. While the EU assigns tailored obligations to each link of the supply chain, Korea operates with two broad categories of development and utilization. The Korea Legislation Research

Institute pointed out a problem with this structure being so broad that even government agencies are included in AI business operators.

Risk classification systems also differ. Korea uses the word "impact." It is a single category of high-impact artificial intelligence. The EU uses the word "risk" and divides it into four levels: unacceptable risk, high-risk, limited risk, and minimal risk. Korea has no category of prohibited practices. The EU outright prohibits practices such as social scoring. Real-time remote biometric identification for law enforcement purposes is prohibited in principle, permitting only narrow exceptions under strict authorization procedures such as searching for missing persons, responding to imminent threats, or identifying suspects of serious crimes. Korea looks at impact; the European Union looks at risk.

While the EU imposes mandatory conformity assessments on high-risk AI, Korea made impact assessments an obligation to endeavor. High-impact duties such as risk management are statutory obligations to perform, but violating them does not immediately incur penalties; administrative fines arise only when reaching the stage of failing to comply with a corrective order. The degree of enforcement is different. The EU raised a sword; Korea offered

guidelines.

The EU has the concept of General-Purpose AI Models (GPAI). Korea does not. Korea defined "generative artificial intelligence," but this is limited to systems that generate content. It is unclear whether general-purpose models that do not generate content—such as classification, reasoning, and tool use—fall under this definition.

7. China, Different Laws for Different Technologies

China walked a completely different path.

It did not enact a single law encompassing all of artificial intelligence. Whenever a technology emerged, it released regulations tailored to that technology. Provisions on the Administration of Algorithm Recommendation in March 2022, Provisions on the Administration of Deep Synthesis in January 2023, and Interim Measures for the Administration of Generative AI Services in August 2023. Three regulations emerged within two years.

The Provisions on the Administration of Algorithm Recommendation divide algorithms into five types: content generation and synthesis, personalized recommendation, selection and ranking, search filter, and dispatch and allocation decision.

Service providers possessing public opinion attributes or social mobilization capabilities are imposed with algorithm filing (备案) obligations. The algorithm type, application domain, and self-assessment results must be filed with the Cyberspace Administration of China within ten working days from the date of service provision.

The Provisions on the Administration of Deep Synthesis list deepfake technologies into six sub-types: text generation and editing, voice synthesis/conversion, non-human image and video generation, human face generation and replacement, 3D reconstruction, and immersive content generation. A concept appears in these provisions that does not exist in Korea or the EU: technical supporter (技术支持者). It establishes persons who do not directly provide services but provide technical support as separate regulated entities. No corresponding definition exists in Korean law, nor in the EU AI Act.

The Interim Measures for the Administration of Generative AI Services define "generative artificial intelligence technology" as "models and related technologies capable of generating content such as text, images, audio, and video." It is similar to Korea's definition of generative artificial intelligence, but while Korea describes even the

mechanism of "imitating the structure and characteristics of input data," China defines it solely by output capability.

China's business operator classification is provider (提供者) and user (使用者). A provider is an organization or individual providing services using generative AI technology, including API provision. A user is an organization or individual generating content using that service. Korea's utilization business operator and China's user (使用者) share similar names, but while Korea's utilization business operator is a person providing products or services externally, China's user is a person generating content. Both scope and obligations differ.

China has no unified definition encompassing all of artificial intelligence. Each regulation defines only the technology under its jurisdiction. Nor does it have extraterritorial provisions. It merely states that it applies when "providing services to the public within the territory of China." It secures jurisdiction through technical blocking rather than explicit extraterritoriality.

What stands out in Chinese regulation is the filing system. Korea aims for ex-post self-regulation. The EU adds prohibited practices to ex-ante conformity assessments. China combines filing with content control. It is a structure where regulatory authorities track through

filing and control ex-post.

8. Sanctions in Three Countries

The teeth of a law lie in sanctions.

The maximum fine under the EU AI Act for violating prohibited practices is up to 35 million EUR or 7 percent of global annual turnover for the preceding financial year, whichever is higher. For global big tech companies, revenue-proportional calculation applies, so actual fines could reach hundreds of billions of KRW.

The Korean AI Basic Act has no administrative fines (fines/penalties based on revenue). Monetary sanctions against a business operator's breach of obligations are limited entirely to administrative fines under Article 43, with a cap of 30 million KRW. In US dollars, approximately 20,000 USD. Between the EU's 35 million EUR and Korea's 30 million KRW, there is a four-digit order-of-magnitude gap.

China has no single fine ceiling spanning all three regulations. The fine set by the Provisions on the Administration of Algorithm Recommendation itself ranges between 10,000 RMB and 100,000 RMB, while the Interim Measures for the Administration of Generative AI Services prioritize warnings and corrective orders, leaving

monetary sanctions to individual laws such as the Cybersecurity Law. However, China possesses sanctions of business suspension and service shutdown. For enterprises seeking to continue operations in the Chinese market, a service shutdown order is heavier than any monetary sanction.

The disparity in sanctions reflects regulatory philosophy.

The EU puts the protection of fundamental rights upfront. It is a preventive approach aiming to block high risks in advance. It sets prohibited practices, imposes mandatory conformity assessments on high-risk AI, and levies up to 7 percent of global turnover for the most severe violations.

Korea placed the promotion of innovation and the protection of rights side by side. Evidence lies in Article 1 of the Act citing "development of artificial intelligence" and "creation of trust base" side by side. There are no prohibited categories, conformity assessment is voluntary, and the ceiling for administrative fines is 30 million KRW. Academia evaluates this law as a framework act. The National Assembly Research Service pointed out that consumer protection is insufficient. The design of this law weighs heavily on innovation, and that is an intentional choice. However, faced with the

question of whether an administrative fine of 30 million KRW can alter the behavior of global AI enterprises, the cost of this choice has not yet been calculated.

China puts social stability and content control upfront. Its vertically divided regulatory structure is a mechanism for swiftly issuing control norms tailored to each new technology as it emerges. Among the three countries, China leads in terms of regulatory speed.

Returning to that radiology specialist at the university hospital in Seoul. This doctor is a user in Korea, a Deployer in Munich, and has no fitting name in Beijing. Because a CT reading assistance system does not generate content, there is no spot that fits neatly into any of China's three technology-specific regulations. Each country attaches different names or no name at all to the same act, placing obligations of different weights upon it. While the AI Basic Act gives room for the industry to breathe, the European Union erects fences of fundamental rights, and China casts nets of filing. Are you subject to this law? Open Article 2, Item 7 of the Act and the answer becomes clear. If you develop and provide artificial intelligence, or provide products or services with it, you are a business operator. If you merely receive and use it, you are a user. And that status attaches on a

service-by-service basis, not an institutional basis. A hospital that is a user in the radiology department becomes a utilization business operator with respect to that chatbot the moment it launches a patient chatbot, and may cross the line toward a development business operator the moment it performs fine-tuning. How thick this law—clothed in the mantle of a basic act—will become amidst the accumulation of subordinate legislation, official interpretations, and precedents is an answer that has not yet emerged.

Part 2. Self-Diagnosis: Which Regulations Apply to Our Service?

Chapter 3. Step 1 Determination: Are We a Generative AI Business Operator?

February 2026, an e-commerce company in Gangnam, Seoul. In the marketing team meeting room, the team leader points to the screen. AI automatically generated product descriptions overnight for 3,000 new items listed on the company's shopping mall. When a photo is uploaded, sentences describing the material, size, and fit appear. A message arrives from the legal team. Under Article 31(2) of the AI Basic Act, outputs generated by generative AI must be labeled. Must "This text was written by AI" be attached to all 3,000 product descriptions? First, the team leader needs to verify one thing: whether the company's service falls under "generative AI" as defined by law.

1. Three Conditions of the Legal Provision

Article 2, Item 5: "An artificial intelligence system that generates text, sound, images, video, or various other outputs by imitating the structure and characteristics of input data."

Deconstructing this definition reveals three conditions.

First, it must "imitate" the structure and characteristics of input data. This is a mechanism where an AI model mimics the patterns and structures of trained data to produce new outputs. A classic example is a large language model learning billions of text data entries and creating new sentences by following their sentence structures and vocabulary patterns.

Second, the type of output must be "text, sound, images, video, or various other outputs." It explicitly specifies text, audio, images, and video, while keeping it open-ended with "various other outputs." Forms such as code, 3D models, and molecular structures can also fall under this category.

Third, it must be an "artificial intelligence system" as defined in Article 2, Item 2. It must be a system possessing autonomy and adaptability that produces outputs through inference. Software that merely repeats pre-programmed rules is not an artificial intelligence system.

All three conditions must be satisfied to constitute generative AI. If even one is missing, it does not qualify.

2. The Key Concept of "Imitation"

Among the three conditions, the one with a broad scope of interpretation is "imitating the structure and characteristics of input data."

Large language models (LLMs) learn billions of text data entries and grasp the probability distribution of tokens. When creating a new sentence, they calculate the probability of the next token based on preceding tokens and output them one by one. While mimicking the sentence structure and vocabulary patterns of the training data, they do not reproduce the training data verbatim. This is imitation.

Generative Adversarial Networks (GANs) learn the statistical distribution of real images and sample new images from that distribution. Diffusion models create new images by learning the process of adding and removing noise from images. Both methods capture the structural characteristics of training data and yield outputs with similar properties.

Imitation is not duplication. Repeating training data as-is constitutes search or database lookup. Imitation is the process of learning the patterns, distributions, and regularities inherent in training data, and creating something unprecedented based on that foundation.

This requirement serves the function of filtering out rule-based systems. Consider a system that creates sentences by merely inserting variables into predetermined templates: "Order No. XXXX for Customer OO has been received." This is not imitating the structure of data; it operates according to rules set by a programmer. It is not generative AI.

3. Clear Cases

ChatGPT is generative AI. It learns from massive volumes of text, imitates the structure and characteristics of text, and generates new text.

The same applies to Midjourney. It learns from massive volumes of images, imitates the structure and characteristics of images, and generates new pictures.

Suno AI involves sound. It learns from massive music audio data, imitates the structure and characteristics of music, and creates new music.

OpenAI's Sora involves video. It learns from massive video data, imitates the structure and characteristics of video, and creates new motion pictures.

GitHub Copilot generates code. It learns from massive source code, imitates the structure and characteristics of code, and generates new programming code. Code falls under "various other outputs" specified by law.

Business operators running these services are generative AI business operators. Obligations of labeling generated outputs and providing prior notifications to users apply to them.

4. Borderline Cases

What is difficult is not the clear cases.

Let us look first at recommendation algorithms. Netflix analyzes viewing history to present "movies you might like." Coupang displays "recommended products" based on purchase records.

Recommendation algorithms select and present existing content. The movie list already exists; the algorithm merely decides which movie to show to whom and does not create new content. Because it selects existing items rather than generating outputs, it does not constitute generative AI.

However, the story changes if the recommendation algorithm also composes personalized phrasing. If AI generates customized copy

such as "If you enjoyed the thriller you watched last week, we also recommend this title," that text generation feature may fall within the scope of generative AI. This is a case where selection and generation are blended into a single service. While the recommendation function itself is not generation, the function that creates recommendation copy can be generation.

Consider search autocomplete. Typing up to "AI B" in a Naver or Google search bar suggests "AI Basic Act." Most autocomplete features display previously entered search queries rather than creating new text. It is difficult to consider this generative AI.

Cases where a search engine integrates AI to present search results as a summary are different. Google's AI Overview and Naver's AI Search Briefing synthesize multiple search results to create entirely new text that did not previously exist. This feature constitutes generative AI. Even within the same search service, autocomplete is categorized as non-generative, while AI summarization is categorized as generative.

Consider translation engines. DeepL translates an English contract into Korean; Papago translates a Japanese thesis into Korean. The output of translation is text. Neural Machine Translation (NMT) learns

from massive parallel corpora and imitates structural characteristics of language to produce text in another language. This process can align with the act of "generating text by imitating the structure and characteristics of input data." However, whether translation should be viewed as "transformation" or "generation" remains without an established interpretation. Translation engines rely on an original text as input, reconstructing its meaning in another language to yield new text. While differing from pure generation in that no output exists without an original text, the underlying mechanism is identical in that a neural network imitates language patterns to assemble sentences. As of July 2026, the Ministry of Science and ICT (MSIT) has not issued an official interpretation regarding this boundary.

Consider AI touch-up filters. When a smartphone camera app takes a photo, AI automatically adjusts brightness, color tone, and sharpness. If it merely adjusts pixel values while maintaining the fundamental structure of the photo, it is difficult to consider that it has "generated" a new image. However, if AI replaces the background entirely or alters a person's appearance, it comes closer to creating an image that did not previously exist. As the degree of touch-up deepens, it crosses the threshold of generation.

The boundary lies not in the name of the technology, but in the nature of the output. Selecting existing items is not generation. Slightly tweaking existing items leaves the generation status ambiguous. Creating something anew that did not exist before is generation.

5. Additional Obligations for Generative Business Operators

Two additional obligations are imposed on business operators that provide generative AI.

First is the obligation to label outputs (Article 31(2)). The fact that an output was created by AI must be indicated on the output. For virtual audio, images, or videos that are difficult to distinguish from reality—so-called deepfake outputs—a notification or label presented in a manner that users can clearly perceive is required (Article 31(3)). The Transparency Guidelines set visible labeling as the general principle for such outputs, making it difficult to satisfy this requirement through invisible watermarks alone. Details of labeling methods are guided by output type in guidelines, while the Enforcement Decree recognizes exceptions when an output is manifestly evident as AI-generated (for instance, outputs generated within an AI service platform) or when outputs are for internal business purposes.

Second is the obligation of prior notification to users (Article 31(1)). When intending to provide products or services utilizing high-impact AI or generative AI, business operators must notify users in advance that the product or service is operated based on such AI. Prior notification is not an obligation unique to generative AI; it also applies to products and services utilizing high-impact AI. However, in the case of generative AI, prior notification is coupled with output labeling, adding an extra layer to the transparency burden.

Both obligations have applied since the enforcement date of January 22, 2026. While a grace period of at least one year is granted for imposing administrative fines, immediate factual investigations are possible in grave matters such as deepfake criminal exploits or loss of human life.

Output labeling is not an end in itself. Its purpose is to enable users to know whether the text they are reading, the sound they are hearing, or the image they are viewing was created by a human or a machine. During the initial 10 days of implementation, 172 inquiries were submitted to the AI Basic Act Support Desk, with transparency-related inquiries accounting for the largest share—a sign that the scope and method of this obligation remain unfamiliar

to business operators.

6. Where the European Union Looks: General-Purpose Functionality

South Korea's AI Basic Act features "generative AI." The European Union's AI Act features "General-Purpose AI (GPAI)." While the two concepts appear similar, their classification criteria differ.

Korea looks at outputs. The criterion is whether it "generates" text, sound, images, or video. Focus is placed on what it creates.

The European Union looks at general-purpose capability. Article 3(63) of the EU AI Act defines a GPAI model as an AI model that "displays significant generality and is capable of competently performing a wide range of distinct tasks." The criterion is not what it creates, but how wide a range of tasks it can perform.

This difference creates a gap.

Suppose there is a general-purpose AI model that performs classification tasks. It determines whether an email is spam, classifies the topic of a document, and identifies objects in an image. This model outputs "classification" results; it does not generate new text, sound, or images. Under EU criteria, it can qualify as GPAI because it

performs multiple tasks. Under Korean criteria, it hardly qualifies as generative AI because it does not "generate" outputs.

AI agents performing reasoning and tool calls encounter the same issue. An agent that manages user schedules, classifies emails, and queries databases may output text, but its core function is calling external tools and making decisions. It is unclear whether such an agent fully falls under generative AI under Korean law. In the EU, it can be encompassed as a system built on top of a general-purpose model (GPAI system).

Korean law lacks a concept corresponding to GPAI. While "generative AI" and "high-performance AI" exist—the latter being a practical shorthand rather than a name explicitly appearing in the statute, referring to AI systems satisfying all three conditions under Article 24 of the Enforcement Decree: cumulative training compute of 10^{26} FLOPs or more, application of cutting-edge AI technology, and potential to exert wide-ranging and grave impacts on life, bodily safety, and fundamental rights—a classification based on "general-purpose functionality" does not exist in the law. The EU has established a granular system distinguishing GPAI models from GPAI systems and imposing additional obligations on GPAI posing systemic

risks. Korea classifies generative AI by a single yardstick of output type, and layers atop it high-performance AI evaluated by compute and risk level. The number of yardsticks differs.

The scope of obligations also differs significantly. Obligations imposed on GPAI providers in the EU have applied since August 2, 2025. These include technical documentation and publicly available summaries of training content. The Code of Practice for GPAI, drafted with the participation of approximately 1,000 stakeholders, was established as a tool for fulfilling these obligations. The three pillars of the Code of Practice are transparency (documentation and 10-year retention), copyright (respecting robots.txt and minimizing infringement), and safety and security (systemic risk management). In Korea, the obligations imposed on generative AI business operators consist of only two: output labeling and prior notification.

Korea imposes labeling obligations on services that yield outputs. The EU builds up obligations based on the general-purpose capability and systemic risk of the model. The regulatory focus differs. Korea concentrates on informing users that "this was created by AI." The EU, as far as general-purpose AI is concerned, concentrates on managing the safety and transparency of the model itself.

7. China: Defining by Capability and Controlling by Registration

China's Interim Measures for the Management of Generative Artificial Intelligence Services (enforced on August 15, 2023) defines "generative artificial intelligence technology" as "models and related technologies capable of generating content such as text, images, audio, and video."

While similar to the Korean definition, there are differences. Korea incorporated the mechanism of "imitating the structure and characteristics of input data" into its definition, looking even at how outputs are produced. China defines it solely by result capability: "capable of generating content," focusing strictly on what can be created.

This difference yields divergent outcomes in borderline cases. Consider a system synthesizing music based on rules. Because its output is sound, it can be encompassed under China's definition. However, because it combines notes according to pre-programmed rules rather than "imitating" the structure of training data, it hardly falls under Korea's definition. Korea's definition is narrower than China's definition.

China's Interim Measures contain unique elements absent in Korea: content control obligations. Service providers must ensure adherence to "Core Socialist Values." Content inciting subversion of state power, terrorism, or ethnic hatred is prohibited. Services possessing "public opinion attributes or social mobilization capabilities" are subject to prior security assessments. Obligations of this nature exist neither in South Korea's AI Basic Act nor in the EU AI Act.

The filing/registration (备案, bei'an) system serves as the primary enforcement tool. Generative AI services with public opinion attributes or social mobilization capabilities must submit a security assessment and register their algorithms with the Cyberspace Administration of China (CAC) within 10 business days after beginning service provision. In practice, most public-facing generative AI services are operated as meeting this requirement. As of November 2025, 611 generative AI services have completed registration, and total algorithm filings reportedly exceed 5,000.

Korea has no such registration requirement. There is no need to file a report with or obtain permission from the government when launching a generative AI service. Only post-launch self-regulation

and transparency obligations of output labeling exist.

China's Measures for Labeling AI-Generated and Synthesized Content, enforced on September 1, 2025, mandated dual labeling: visible labels must be attached simultaneously with metadata or digital watermarks. Mandatory national technical standards (GB 45438-2025) were even enacted. In contrast, Korea's output labeling obligation operates by having the Enforcement Decree lay down the broad framework of labeling methods while guidelines instruct details, with no mandatory technical standards.

Comparing the approaches of the three jurisdictions in a single sentence: Korea captures AI by what it creates (output type); the EU captures AI by how widely it can be used (general-purpose functionality); China captures AI by whether it possesses the capability to create (result capability).

8. The 3,000 Product Descriptions

Returning to the e-commerce company in Gangnam.

The AI system used by this company takes product photos as input and generates text. It learned from a massive volume of product description data and creates new descriptions by imitating its

structure and characteristics. It satisfies all three conditions of Article 2, Item 5. It is generative AI.

If this company built the AI model directly, it is a developer business operator. If it calls external APIs and provides services to users with the resulting outputs, it is a user business operator. In either case, as a business operator, obligations of output labeling and prior notification apply.

Must "Written by AI" be attached to all 3,000 product descriptions? An exception is recognized when an output is manifestly evident as AI-generated. However, if users cannot distinguish between human-written product descriptions and AI-written product descriptions, labeling is required.

The same team leader opens another screen, showing the "Recommended Products" list of the company's shopping mall. AI analyzes purchase history for each customer to select and present a list. This is selection, not generation. Additional obligations of generative AI do not apply. Within the same company on the same day, the same AI technology is generative on one side and non-generative on the other.

There are places where the law is clear, and places where it is not. ChatGPT, Midjourney, and Suno AI are clear. Recommendation algorithms are clearly outside. Translation engines, search autocomplete, and AI touch-up filters stand on the boundary line. Guidelines from the Ministry of Science and ICT and consultations at the AI Basic Act Support Desk are filling in this boundary line by line, but technology moves faster than legal interpretation. How far the perimeter of the category of generative AI will extend will be revealed as services multiply and precedents accumulate.

Chapter 4. Two-Step Determination: Are We a High-Impact AI Business Operator?

It is the autumn of 2025 in a meeting room of a mid-sized HR management company. This company operates an AI interview analysis system. When an applicant undergoes a video interview, the AI analyzes voice tone, answer content, and changes in facial expression to generate a suitability score. The HR team of a client conglomerate uses this score as a reference for deciding whether applicants pass the first-round interview. On January 22, 2026, the AI Basic Act comes into effect. The legal officer of this company opens Article 2, Subparagraph 4 of the Act. Subparagraph (g): "Determinations or evaluations that have a significant impact on individuals' rights and obligations, such as recruitment and loan screening." Their eyes stop there. Does our system fall under this? If it does, the company must establish risk management measures, explain the basis of determination, prepare user protection measures, maintain a human-in-the-loop system, and retain documents for five years. If it does not, those obligations do not apply. Hundreds of millions of won in compliance costs hinge on this single line of judgment.

1. What Is High-Impact Artificial Intelligence?

Article 2, Subparagraph 4 of the AI Basic Act defines it as: "An artificial intelligence system that has a significant impact on, or poses a potential risk to, human life, physical safety, and fundamental rights, and is utilized in any of the areas specified in the following items."

Two conditions must be met simultaneously. It must be used in an area enumerated by law, and it must have a significant impact on human life, physical safety, or fundamental rights. Even if it falls under an enumerated area, it is not high-impact if the impact is not significant. Even if the impact is significant, it is not high-impact if used outside the enumerated areas.

This is the two-step determination structure.

Step 1. Domain applicability. In which of the eleven areas listed from items (a) to (k) of Article 2, Subparagraph 4 of the Act is this AI system utilized?

Step 2. Determination of significance. Does this system, within that domain, have a significant impact on, or pose a potential risk to, human life, physical safety, and fundamental rights?

Article 25 of the Enforcement Decree also sets out the procedure for this determination. A business operator must examine in advance whether its system qualifies as high-impact and, if necessary, may request confirmation from the Minister of Science and ICT. A default period of 30 days is granted for the Minister's response upon receiving the request. A 30-day extension is allowed once. For business operators, it is safer to document the grounds for their preliminary examination.

2. The Eleven Areas

Article 2, Subparagraph 4 of the Act lists eleven areas from items (a) to (k).

Item (a). Supply of energy pursuant to the Energy Act. This applies when AI is used in power supply grids, gas pipeline control, or heat supply systems. If a power supply-and-demand forecasting model malfunctions, a wide-area blackout occurs. The rolling blackout across South Korea on September 15, 2011, was caused by an error in power supply-and-demand forecasting. In an era where AI takes over such forecasting, the failure of that AI is directly the failure of infrastructure.

Item (b). Production process of drinking water pursuant to the Drinking Water Management Act. This targets AI that analyzes water quality data and determines chemical dosage at water treatment plants. Incorrect chemical dosage leads to water pollution, and water pollution leads to disease. This area has a narrow scope. It is limited to the "production process." It does not include tap water pipe management or drinking water distribution.

Item (c). Establishment and operation of the healthcare provision and utilization system pursuant to the Framework Act on Health and Medical Services. This includes hospital information systems, telemedicine platforms, patient allocation algorithms, and medical resource allocation systems. The scope is broad, covering the overall establishment and operation of the healthcare system. Systems where AI assists in emergency patient triage, systems where AI decides bed allocation, and systems where AI analyzes health checkup results to recommend additional tests can all fall under item (c).

Item (d). Development and utilization of medical devices pursuant to the Medical Devices Act and digital medical devices pursuant to the Digital Medical Products Act. AI image reading software, AI-based ECG

analyzers, and AI diagnostic assistance programs fall under this. The portion of this item regarding digital medical devices took effect on January 24, 2026—two days later than the original Act's enforcement date of January 22. This is because the Digital Medical Products Act came into effect on the same day. The portion concerning medical devices under the Medical Devices Act applied from January 22 along with the original Act. Medical device business operators already have a licensing and approval framework under the Medical Devices Act. The AI Basic Act imposes high-impact AI obligations on top of that. However, there is a provision for easing redundant regulations, allowing business operators that have already fulfilled their obligations under the Medical Devices Act to be recognized as having fulfilled the obligations under the Basic Act within that corresponding scope.

Healthcare under item (c) and medical devices under item (d) have overlapping areas. A single AI image reading system could fall under both items, as it is a medical device while simultaneously being used in the operation of the healthcare system. In this case, since it qualifies as high-impact under either item, there is little practical difference in the scope of obligations.

Item (e). Safety management and operation of nuclear materials and nuclear facilities pursuant to the Act on Physical Protection and Radiological Emergency Preparedness for Nuclear Facilities, etc. This applies when AI is applied to nuclear power plant control systems, nuclear fuel management, or radiation monitoring systems. Nuclear energy is a domain where the cost of failure is virtually infinite. Business operators using AI in this area are already subject to strict regulations under the Nuclear Safety Act framework.

Item (f). Analysis and utilization of biometric information for crime investigation or arrest. The Act lists face, fingerprint, iris, and palm vein as examples of biometric information. Facial recognition AI, fingerprint matching AI, and iris scan analysis AI used for crime investigation purposes are covered. The purpose is restricted: it must be "for crime investigation or arrest." Facial recognition for immigration control or building security does not fall under item (f). However, there is potential for it to fall under item (g) or item (i).

The EU AI Act handles biometric identification more strictly than South Korea. The use of real-time remote biometric identification systems for law enforcement purposes is prohibited in principle, allowing only narrow exceptions such as terror prevention or

searching for missing children. Korean law contains no such outright prohibition. It merely classifies biometric AI for crime investigation as high-impact and imposes safety obligations, without blocking its use.

Item (g). Determinations or evaluations that have a significant impact on individuals' rights and obligations, such as recruitment and loan screening. The word "such as" (or "etc.") leaves the scope open. Recruitment and loan screening are examples. Insurance underwriting, student loan evaluation, lease eligibility determination, and social security benefit eligibility adjudication can all be subsumed under "determinations that have a significant impact on individuals' rights and obligations."

The AI interview analysis system mentioned earlier is highly likely to fall under item (g). If this system's suitability score substantially affects the hiring decision, the significance requirement of the Step 2 determination is also met. However, if this score serves merely as reference material and the HR manager makes the decision independently, there is room for it to be excluded from the significance determination. Ultimately, the weight that AI carries in the recruitment decision determines whether it qualifies as high-impact.

Item (h). Major operation and management of transportation means, transportation facilities, and transportation systems pursuant to the Transportation Safety Act. Driving control of autonomous vehicles, flight control assistance, railway operation management, and signal system optimization are targeted. A qualifying modifier is attached: "major operation and management." Route recommendations in traffic information apps or navigation systems are not "major operations" of transportation facilities, making it highly likely they do not fall under item (h). However, AI that controls steering and braking in autonomous vehicles is the very major operation of a means of transportation.

Item (i). Decision-making by the State, local governments, and public institutions pursuant to Article 4 of the Act on the Management of Public Institutions that affects citizens, such as qualification verification and determination or fee collection necessary for providing public services. This applies when AI is used for grading civil service exams, adjudicating welfare benefit eligibility, or setting public utility rates. Since the subjects are the State, local governments, and public institutions, private business operators are rarely the direct parties. However, when a private company supplies an AI system to a public institution and that system is used for

decisions falling under item (i), the company may bear high-impact obligations as a developer business operator.

Item (j). Student evaluation in early childhood education, primary education, and secondary education pursuant to the Framework Act on Education. Systems where AI grades student exams, evaluates academic achievement, or diagnoses learning levels are targeted. The scope is limited. Higher education is excluded. It applies only from early childhood to high school students. Learning aids (such as AI tutors or problem recommendations) that are not student evaluations do not constitute evaluations and are thus not subject to item (j).

Item (k). Other areas that have a significant impact on human life, physical safety, and fundamental rights protection as prescribed by Presidential Decree. As of July 2026, no additional areas have been designated by Presidential Decree. This item is an open-ended clause, left in place to allow new domains to be added in response to technological advances or societal changes.

3. Practice of the Two-Step Determination

The Enforcement Decree and the High-Impact AI Determination Guidelines operate the two-step determination as follows.

Step 1 is relatively clear. Cross-check where our system is used among the 11 areas: energy supply, drinking water production, healthcare, medical devices, nuclear power, biometric crime investigation, determination of rights and obligations such as recruitment or loan screening, transportation, public services, and student evaluation. If no relevant area applies, it is not high-impact. If a relevant area applies, proceed to Step 2.

Step 2 is tricky. What constitutes a "significant impact"? The law states: "having a significant impact on, or posing a potential risk to, human life, physical safety, and fundamental rights." Article 25, Paragraph 2 of the Enforcement Decree provides factors to consider in making this determination: the impact, severity, and frequency of risks that could be posed to human life, physical safety, and fundamental rights, as well as the specific characteristics of each domain of application.

Let us break this down into specific cases.

An AI recruitment interview system. Step 1 determination: It falls under "recruitment" in item (g). Step 2 determination: If the AI's suitability score effectively determines whether an applicant passes the first-round interview, it significantly impacts the applicant's

fundamental right to freedom of choice of occupation. This system is high-impact. Conversely, even with the same system, if the structure is such that the AI score is merely for reference and HR personnel independently interview all candidates to make the final decision, it is difficult to view the AI's impact as significant. In a structure where "a human decides and AI merely assists," it may not pass Step 2. This boundary is where frequent disputes will arise in practice.

An AI credit scoring system. Step 1 determination: It falls under "loan screening" in item (g). Step 2 determination: If the credit score is directly linked to loan approval or interest rate determination, it significantly impacts an individual's property rights. If it is a credit rating company already subject to the Credit Information Act, under the provisions for easing redundant regulations, it may be recognized as having fulfilled the obligations under the Basic Act to the extent it has fulfilled its obligations under the Credit Information Act.

An AI student evaluation system. Step 1 determination: It falls under "student evaluation" in item (j). Step 2 determination: If this AI grades midterm or final exams and those grading results are reflected in academic records (GPA), it significantly impacts the student's right to receive education and opportunities for higher education. It is

high-impact. If the AI merely recommends practice problems or displays learning progress, it is not an "evaluation," so it does not fall under item (j), and there is no need to proceed to Step 2.

4. Is It "Impact" or "Risk"?

Senior Research Fellow Won-Jun Jung of the Korea Legislation Research Institute pointed out the terminological confusion of this Act in a September 2025 report.

Let us revisit the definition in Article 2, Subparagraph 4 of the Act: "An artificial intelligence system that has a significant impact on, or poses a potential risk to, human life, physical safety, and fundamental rights." Both "impact" and "risk" appear within a single sentence. "Has a significant impact on" is followed by "poses a potential risk to."

"Impact" and "risk" are distinct concepts. Impact refers to the magnitude of a consequence that has already occurred or is expected to occur. Risk is the product of the probability and severity of a harmful outcome occurring. The title of the Korean law uses "high-impact." Yet, the word "risk" is also included in the defining statutory provision.

Research Fellow Won-Jun Jung's critique is as follows: as "impact" and "risk" are used side-by-side within the law, confusion arises for business operators as to what standard they should base their judgment on. Furthermore, in the phrase "human life, physical safety, and fundamental rights," the word "and" (∧) connects safety and fundamental rights in parallel, leading to conflicting interpretations over whether a system qualifies if it affects fundamental rights even without affecting safety, or whether it must affect both safety and fundamental rights.

From a business operator's perspective, it translates to this: an AI recruitment interview system does not affect human life or physical safety. It does affect fundamental rights (freedom of choice of occupation). If "and" (∧) is read conjunctively ("and"), an interpretation can emerge that as long as there is no impact on life and physical safety, it is not high-impact. If "and" (∧) is read disjunctively ("or"), an interpretation arises that impacting fundamental rights alone makes it high-impact. The guidelines of the Ministry of Science and ICT operate in the latter direction—that even cases affecting only fundamental rights can qualify as high-impact.

Looking back at the deliberation process in the National Assembly's Science, ICT, Future Planning, Broadcasting and Communications Committee, choosing "impact" instead of "risk" in the Act's title was a political compromise to avoid placing a negative stigma on technology. As shown by Lawmaker Hae-Min Lee's remarks, there was a judgment that the word "impact" was more neutral and comprehensive than "risk." However, by using both impact and risk simultaneously in the statutory definition, tension arose between the neutrality of the title and the clarity of the definition.

The European Union does not have this issue. From the start, it used the word "risk" and built a framework across four tiers: unacceptable risk, high-risk, limited risk, and minimal risk. Terminology and structure align seamlessly.

5. High-Risk AI in the European Union: Dual Pathways and Derogations

The EU AI Act classifies high-risk AI into two pathways.

Pathway 1: The product safety pathway (Article 6(1)). Under EU harmonization legislation (product-related directives listed in Annex I, such as medical devices, machinery, toys, aviation, and motor vehicles), if AI is used as a safety component of a product subject to third-party conformity assessment, it is high-risk. Diagnostic AI

software embedded in medical devices is a prime example. In this pathway, existing product safety frameworks overlap with AI regulation.

Pathway 2: The Annex III pathway (Article 6(2)). This targets stand-alone AI systems used in the areas listed in Annex III: biometrics, critical infrastructure, education and vocational training, employment and workers management, access to essential private and public services, law enforcement, migration, asylum and border control management, and administration of justice and democratic processes. While it overlaps considerably with Korea's 11 domain areas, there are also differences. Whereas the EU explicitly specifies migration and border control, and administration of justice and democratic processes, South Korea has not listed these areas. While Korea separated energy supply, drinking water production, and nuclear energy into distinct items, the EU grouped them under a single heading of "critical infrastructure."

The EU has a mechanism that Korea lacks: a derogation mechanism (Article 6(3)). Even for AI systems in areas listed under Annex III, if they do not pose a significant risk to human health, safety, or fundamental rights, they are not classified as high-risk. Performing a narrow

procedural task, aiming to improve the result of a previously completed human assessment, or assisting in decision-making without replacing human assessment serve as grounds for derogation. However, profiling of natural persons is always high-risk; no derogations apply.

Korean law lacks this derogation mechanism. If it falls under the 11 areas and has a significant impact, it is high-impact. Unlike the EU, there is no statutory path to exit by claiming, "This system is used in an Annex III area, but since its practical risk is low, it is not high-risk." While a similar conclusion might be reached in Step 2 through the interpretation of "significant impact," explicit derogation conditions are not enumerated as in the EU.

Another difference between South Korea and the EU lies in the intensity of obligations. The EU imposes mandatory conformity assessment on high-risk AI. Prior to placement on the market, technical documentation must be drawn up, a quality management system established, an EU declaration of conformity issued, and CE marking affixed. Cases requiring third-party assessment also exist. In non-compliance, administrative fines up to EUR 15 million or 3 percent of global annual turnover may be imposed.

The obligations of high-impact AI business operators in Korea are different in nature. Article 34, Paragraph 1 of the Act specifies them in five subparagraphs within a single article. Business operators must establish and operate risk management measures. They must establish and implement measures to explain the key criteria for determination and an overview of training data. They must establish and operate user protection measures. They must maintain a human management and supervision system. They must prepare and retain documentation evidencing compliance for five years. Added to this is the disclosure obligation under Article 27, Paragraph 1 of the Enforcement Decree: key contents of risk management measures, key contents of explanation measures, user protection measures, and the name and contact details of the manager/supervisor must be posted at offices, business places, or websites, though trade secrets may be omitted. While the list of obligations is substantive, there are no prior entry barriers like third-party conformity assessments. Nor is the structure such that an administrative fine is immediately imposed for violating these duties. Only when an operator fails to comply after receiving a corrective order from the government is an administrative fine of up to KRW 30 million imposed. The government announced a policy to operate a grace period of at least one year after enforcement.

6. China: Dividing Domains Without Risk Classification

China does not use the categories of "high-risk AI" or "high-impact AI" at all.

Unlike the EU, China has no horizontal classification scheme traversing all AI by risk tier. It merely maintains principles of classification and graded management within individual regulations. Instead, regulations are divided by technology type: the Provisions on the Management of Algorithmic Recommendations for Internet Information Services (enforced March 2022) for algorithmic recommendations; the Provisions on the Administration of Deep Synthesis of Internet Information Services (enforced January 2023) for deepfakes; and the Interim Measures for the Management of Generative Artificial Intelligence Services (enforced August 2023) for generative AI. In 2025, measures on labeling AI-generated content were added; in July 2026, interim measures for anthropomorphic AI interaction services were enforced; and prior to that in May, implementation opinions regarding norms on AI agents were announced.

This is the inverse structure of Korea, which encompasses all AI under a single law and divides them within it into high-impact and general AI. China released regulations tailored to each technology as

it emerged. Each regulation defines only the technology within its jurisdiction and imposes obligations tailored to that technology.

In China, the closest concept to risk classification is the criterion of "public opinion attributes or social mobilization capacity." The Interim Measures for Generative AI Services require security assessments for services possessing such attributes. AI services with "social mobilization capacity" must apply to the Cyberspace Administration of China (CAC) for a security assessment in advance, and can only provide services upon passing this assessment. This is a standard different from Korea's high-impact or the EU's high-risk. The axis of evaluation is the maintenance of social stability, not concerns over fundamental rights infringement.

China is preparing a comprehensive artificial intelligence law. In its annual legislative work plan in May 2026, the State Council stated that it would "accelerate comprehensive legislation for the healthy development of artificial intelligence." Scholarly proposals (jointly submitted by academia, including China University of Political Science and Law, in March 2024) include a permit management system based on an AI negative list, with academic forecasts envisioning passage by the Standing Committee of the National

People's Congress around late 2027 and enforcement in 2028. Once this comprehensive law is enacted, China is also expected to introduce risk tier classification in some form.

Comparing the three jurisdictions, the picture is as follows:

South Korea established a single category of "high-impact," enumerated 11 areas, and determines applicability through a two-step evaluation. There are no prohibited categories.

The EU set up a four-tier risk classification, established prohibited practices, created dual pathways for high-risk AI, and even incorporated a derogation mechanism. Its framework is highly sophisticated.

China, lacking EU-style risk tiering, divides regulations by technology type and combines content control with prior filing/registration.

7. The Business Operator's Determination Procedure

Here is a summary of the procedure by which a business operator self-assesses whether its system qualifies as high-impact artificial intelligence.

Step 1. Is our system an artificial intelligence system? Does it satisfy the requirements of Article 2, Subparagraph 2 of the Act (autonomy, adaptiveness, reasoning)? Simple automation that merely follows preset rules is interpreted as failing to meet these requirements.

Step 2. Are we an artificial intelligence business operator? Do we develop and provide this system, or do we receive it and create a service to provide? If we merely receive and use it, we are a user.

Step 3. Is it used in one or more of the 11 areas? If not applicable, it is not high-impact.

Step 4. Does it have a significant impact on life, physical safety, or fundamental rights? Consider the severity of the impact and risk, frequency of occurrence, and domain-specific characteristics of application.

Step 5. Document the basis for the determination. Whether it qualifies as high-impact or not, it is safer to record the evaluation process and grounds. If self-assessment is difficult, a business operator may request confirmation from the Minister of Science and ICT pursuant to the procedure under Article 25 of the Enforcement Decree, and upon receiving the request, the Minister will issue a

response within 30 days (up to 60 days with a single extension). If there is an objection to the response, a request for re-confirmation may be made within 10 days from the date of receiving the response.

If it qualifies as high-impact, the obligations specified in five subparagraphs under Article 34, Paragraph 1 of the Act will follow: establishment and operation of risk management measures; establishment and implementation of measures to explain the basis of determination to users; establishment and operation of user protection measures; human management and supervision; and the creation and five-year retention of compliance documentation.

8. Systems on the Boundary

There are systems where qualification as high-impact is unclear.

An AI-based insurance underwriting system. AI analyzes insurance applications to determine approval, rejection, or premium rates. This could fall under "determinations that have a significant impact on individuals' rights and obligations, such as loan screening" in item (g). The key is whether "such as" includes insurance underwriting. Rejection of insurance coverage directly affects an individual's economic rights. There is a possibility it could be classified as high-impact.

An AI-based traffic information service. It analyzes real-time traffic data to recommend optimal routes. Item (h) states "major operation and management of transportation means, transportation facilities, and transportation systems." Route recommendation is not a major operation or management of a transportation system. It is an auxiliary service providing convenience to drivers. It is highly likely not to be high-impact.

An AI-based apartment management system. It optimizes energy consumption to manage communal heating and electricity. Item (a) states "supply of energy pursuant to the Energy Act." Does energy management inside an apartment complex qualify as energy supply under the Energy Act? Energy supply under the Energy Act refers to the act of power or gas companies supplying energy to consumers. An apartment management system allocates supplied energy rather than supplying energy itself. It is highly likely not to fall under this item.

An AI-based national scholarship eligibility determination system. It falls under "qualification examination, verification, and determination by the State, local governments, and public institutions" in item (i). In the Step 2 determination as well, it directly

impacts students' educational opportunities. It is high-impact.

As borderline cases accumulate, the indeterminacy inherent in the phrase "significant impact" of the Step 2 determination becomes evident. This indeterminacy is a feature of design, not a flaw. Since the Basic Act cannot classify all cases in advance, the framework leaves room for interpretation, to be filled in through guidelines, authoritative interpretations, and accumulated precedents. For business operators, however, this indeterminacy represents a cost. An incorrect assessment results in a breach of obligation, while over-compliance leads to unnecessary expenditure.

9. Looking at the Frameworks of the Three Jurisdictions

Returning to the legal officer of the HR management company.

This AI interview analysis system is highly likely to qualify as high-impact under Korean law. Step 1: recruitment domain under item (g). Step 2: significant impact on the applicant's fundamental right to freedom of choice of occupation. The company must establish risk management measures, create a system to explain the basis of determination, prepare user protection measures, equip human intervention mechanisms, and retain documentation for five years.

If the same system were launched in the European market, it would fall under the employment domain of Annex III of the EU AI Act and be classified as high-risk AI unless it meets derogation criteria. It would need to undergo conformity assessment, prepare technical documentation, and bear CE marking. However, with the 2026 Digital Omnibus amendments, the application of Annex III high-risk AI obligations has been postponed to December 2, 2027. Comparing sanction ceilings, EUR 15 million alone is more than 700 times South Korea's maximum administrative fine of KRW 30 million, and for conglomerates to which the 3 percent of revenue standard applies, the gap widens several thousand-fold.

If the same system were provided as a service in China, as there is no category of high-risk as such, one would have to check regulation by regulation for applicable specific rules. Since it is not generative AI, the Interim Measures for Generative AI do not apply. If algorithmic recommendation is used, the filing/registration obligations under the Provisions on the Management of Algorithmic Recommendations may apply.

Behind the single-line definition of high-impact artificial intelligence, business operator costs and user protections are

inextricably intertwined. Although the Act is titled "high-impact," the word "risk" is also present in its statutory text. The significance standard of the Step 2 determination is still being filled in through real-world cases. As confirmation procedures by the Minister of Science and ICT, guideline operations, and consultation cases at the AI Basic Act Support Desk accumulate one by one, the contours of the boundaries will take shape. What business operators can do at this point is to follow the five-step determination procedure, document the grounds for their judgment, and judge conservatively if on the borderline. Rather than receiving an inquiry and having it revealed that a system was not high-impact after assuming it was not, assuming it is high-impact, fulfilling the obligations, and later finding out it was excessive carries far less risk.

Chapter 5. Three-Stage Assessment: Are We High-Performance, Cutting-Edge AI Business Operators?

In 2023, the AI research organization Epoch AI estimated the training compute of GPT-4: approximately 2×10^{25} power FLOPs. This is the result of thousands of GPUs running non-stop for several months. The estimated training cost was 100 million USD, or approximately 140 billion KRW. Even this GPT-4 may not reach the high-performance AI criteria set by the Enforcement Decree of South Korea's AI Basic Act. The compute threshold set in Article 24 of the Enforcement Decree is a cumulative training compute of 10^{26} power FLOPs. This is five times the estimate for GPT-4.

1. The Threshold of 10^{26} Power

The concept of high-performance artificial intelligence does not appear among the twelve definitions in Article 2 of the AI Basic Act. Article 32 of the Basic Act declares that obligations to ensure safety will be imposed on artificial intelligence systems that exceed criteria prescribed by Presidential Decree, and Article 24 of the Enforcement Decree specifies those criteria. There are three criteria, and all three must be met: the cumulative compute used for training must be 10^{26} or more; the system must be configured and

operated by applying the most cutting-edge artificial intelligence technology among currently utilized technologies; and there must be concern that the risk level could have a broad and significant impact on human life, physical safety, or fundamental rights. Even if an AI exceeds the compute threshold, if it fails to satisfy the remaining two requirements, it does not fall into this category.

FLOP is an abbreviation for floating-point operation. It is a unit representing a computer calculating numbers with floating points once. To train an AI model, data is repeatedly read, errors are calculated, and weights are adjusted billions of times. The sum of all those calculations is the cumulative training compute.

Ten to the 26th power is a number with twenty-six zeros following a 1. To generate this number, tens of thousands of state-of-the-art GPUs must run non-stop for several months. Adding power, cooling, and data center construction brings the investment to hundreds of billions of KRW. Fewer than ten organizations worldwide can sustain a scale like this.

Article 24, Paragraph 2 of the Enforcement Decree delegates the specific calculation method for cumulative compute to be determined by a public notice of the Minister of Science and ICT. In

early 2026, the Ministry of Science and ICT issued an advance administrative notice for the draft public notice on the implementation method of safety measures. What the public notice can determine is limited to the method of counting compute. Changing the target criteria itself—which artificial intelligence systems are targeted—requires amending the Enforcement Decree.

2. Which Models Cross This Threshold?

Faced with the number 10 to the 26th power, we must examine where actual models stand.

The problem is that hardly anyone knows the exact training compute. None of OpenAI, Google DeepMind, Anthropic, or Meta has officially disclosed the training compute of their models.

Independent research organizations like Epoch AI estimate it by combining publicly available information, technical papers, and industry testimony. These are estimates, not official figures.

According to Epoch AI's estimates, the training compute of GPT-4 is approximately 2 times 10 to the 25th power FLOPs. Google's Gemini 1.0 Ultra is estimated at approximately 5 times 10 to the 25th power. Meta's Llama 3.1 405B is approximately 3.8 times 10 to the 25th power. All are in the 10 to the 25th power range. They fall short of 10

to the 26th power.

However, this landscape is already moving. Epoch AI estimates that the training compute of some frontier models trained after 2024, such as xAI's Grok 3, exceeds 10 to the 26th power. This remains strictly an independent organization's estimate, with no official figures confirmed by the developer.

Korea's threshold is set at a height that even the world's top-tier models of 2023 could not cross. This is no accident.

3. A Ten-Fold Gap Between Korea and the European Union

The EU AI Act uses 10 to the 25th power FLOPs as the criterion for identifying general-purpose AI models (GPAI) with systemic risk. This is a ten-fold difference from Korea's 10 to the 26th power. It is not merely a difference in numbers. The EU's 10 to the 25th power is a mechanism that presumes systemic risk once that line is crossed, whereas Korea's 10 to the 26th power is a combined threshold that applies only after two additional requirements—the application of cutting-edge technology and concerns of significant risk—are satisfied.

Ten-fold is not a numerical difference, but a policy difference. The scope of regulatory targets changes.

Under the EU's 10 to the 25th power criterion, GPT-4 (approx. 2 times 10 to the 25th power), Gemini Ultra (approx. 5 times 10 to the 25th power), and Llama 3.1 405B (approx. 3.8 times 10 to the 25th power) all exceed the threshold. OpenAI, Google, and Meta, which developed these models, are classified under EU law as providers of GPAIs with systemic risk and bear additional obligations such as transparency obligations, technical documentation, and systemic risk mitigation measures.

Under Korea's 10 to the 26th power criterion, the same models fall below the threshold. GPT-4 is one-fifth of the criterion. Gemini Ultra is half. Under Korean law, these models are not high-performance artificial intelligence.

The breadth of capturing regulatory targets differs. The EU set a broad perimeter for potential risks and established a broad regulatory target. Korea took a stance of capturing only top-tier models that can actually exert a broad impact.

The intention is clear. It is a decision to lighten the burden on business operators by narrowing regulatory targets while the domestic AI industry is growing. The high threshold of 10 to the 26th power is a mechanism ensuring that, at present, virtually no domestic business operator needs to fulfill the high-performance AI obligations of this law.

There are counterarguments to this decision. In academia, points are raised that Korea's criterion is an order of magnitude higher than the EU's, potentially creating regulatory blind spots. Training compute is surging every year, so today's high threshold may become ordinary tomorrow. If the criterion were fixed in statutory law, it would be a serious issue, but Korean law delegated the criterion to the Enforcement Decree. Amending the Presidential Decree can lower the threshold.

China has not adopted compute-based classification. Because it controls technology through individual regulations by sector, there is no structure dividing regulatory levels based on training compute. However, in the process of discussing the enactment of a comprehensive AI law, a proposal to impose mandatory safety impact assessments on models exceeding 100 billion parameters is being

raised. Using the number of parameters rather than compute (FLOPs) as a yardstick is an approach that differs from Korea and the EU.

4. What Is Required of High-Performance AI Business Operators?

Business operators that develop and provide artificial intelligence satisfying all three requirements of Article 24 of the Enforcement Decree are subject to obligations to ensure safety. Article 32 of the Basic Act establishes the ground for the obligation, and Article 24 of the Enforcement Decree restricts the targets.

The scope of obligations covers the entire lifecycle of the model.

Risks must be identified. Operators must identify in advance what harmful results the model could produce and what possibilities for misuse exist. The identified risks must be assessed. The likelihood of occurrence and the scale of damage must be estimated. Mitigation measures must be taken according to the assessment results.

Technical supplements, usage restrictions, and monitoring systems are the methods. This process and the results must be documented and submitted.

A risk management system to monitor and respond to AI-related safety incidents must be established. This obligation continues from

the point of model training to the point of service termination.

There are overlapping areas with the obligations of high-impact AI business operators. Under Article 34, Paragraph 1, high-impact AI business operators bear five action obligations: establishing risk management measures, explaining the basis for decisions, protecting users, human management and supervision, and drafting and retaining documentation for five years. For high-performance AI business operators, obligations to ensure safety are added on top of those. If a model exceeding 10^{26} FLOPs is used in high-impact domains such as hiring decisions or loan screening, the obligations of both sides must be fulfilled.

The EU AI Act imposes more specific obligations on GPAI models with systemic risk. Model evaluations and red-teaming tests must be conducted, systemic risks must be tracked and mitigated, and serious incidents must be reported to the EU AI Office. Obligations to draft and maintain technical documentation and publish a sufficiently detailed summary of training data are borne jointly by all GPAI providers regardless of systemic risk status. The Code of Practice for GPAI, assessed as suitable by the European Commission in August 2025, is built around three pillars: transparency, copyright, and

safety. It requires specific responses to four types of systemic risk: CBRN (chemical, biological, radiological, and nuclear) threat capabilities, loss of control, cyberattack assistance, and harmful manipulation and deception.

Obligations to ensure safety under Korean law are mere outlines compared to this. While there is a framework of risk identification, assessment, mitigation, and documentation, detailed implementation methods such as red-teaming tests or systemic risk reporting are in the process of being filled in through public notices and guidelines.

5. The Distance Between Domestic AI Companies and 10 to the 26th Power
Naver's HyperCLOVA X. Samsung Electronics' Gauss. LG AI Research's EXAONE. KT's Mi:dm. Kakao's Kanana.

As of 2026, there are several companies in Korea developing proprietary large language models. The training compute of these models has never been officially announced, but according to general industry estimates, it remains in the range of 10 to the 23rd power to 10 to the 24th power FLOPs. This is a gap of one hundred to one thousand times compared to 10 to the 26th power.

GPU supply and demand is the core bottleneck. Securing tens of thousands of NVIDIA A100 or H100 GPUs for an extended period to run large-scale training requires investments of trillions of KRW. OpenAI received investments on the scale of tens of trillions of KRW from Microsoft. Google operates its own TPU infrastructure. Meta deployed hundreds of thousands of GPUs in its data centers. Domestic companies are in a different dimension from this scale. The National AI Strategy Committee's deliberation and resolution of the Implementation Plan for the National AI Computing Center in September 2025 is also a government-level response to this gap.

To the extent confirmed by publicly available data, there appear to be no domestic AI business operators at present that satisfy the high-performance AI criteria of Article 24 of the Enforcement Decree. Although the law has been enacted, the regulated entities targeted by this provision are not Korean companies.

Business operators that can become regulated entities are global Big Tech. If frontier models developed by OpenAI, Google, Meta, or Anthropic are provided to the Korean market while satisfying the three requirements of Article 24 of the Enforcement Decree, the Basic Act applies. Under the extraterritorial application clause of Article 4,

Paragraph 1, even acts performed overseas must comply with this law if they affect the domestic market or users. If the relevant foreign business operator falls under the criteria for designating a domestic representative (one of four criteria: previous year's revenue of 1 trillion KRW or more, previous year's revenue in the AI service sector of 10 billion KRW or more, daily average of 1 million or more domestic users during the immediately preceding 3 months as of the end of the previous year, or a history of administrative fines imposed under this law), it must designate a domestic representative and handle duties such as submitting the results of fulfilling safety measures through that representative.

As far as high-performance AI is concerned, Korean law targets foreign companies, not Korean companies. Domestic companies fall short of this criterion, and the likelihood of reaching it in the near future is also low.

6. The Position of Business Operators Using APIs

Suppose a domestic startup calls OpenAI's GPT-5 API to create a legal document review service. If GPT-5 is a high-performance AI exceeding 10^{26} FLOPs, does this startup also become a high-performance AI business operator?

No.

This startup did not develop GPT-5. It uses it. Under Article 2, Item 7, Sub-item (b), it is a user business operator. The obligation to ensure safety for high-performance AI belongs to the business operator that develops and provides that AI—here, OpenAI. The act of calling an API is unrelated to training compute. At the point of the API call, the training of the model is already completed.

Whether this startup bears high-impact AI obligations is a separate question. If legal document review falls under a high-impact domain (Article 2, Item 4, Sub-item (g): decisions having a significant impact on individual rights and duties), this startup, as a user business operator, bears obligations related to high-impact AI. It must establish risk management measures, explain the basis for decisions, and maintain a system where human intervention is possible. That is an obligation based on the domain of application, unrelated to 10 to the 26th power FLOPs.

Business operators that do not directly develop large language models but use them via APIs are not high-performance AI business operators. However, depending on which domain and in what manner they use the model, they may bear the obligations of a

high-impact AI business operator. High-performance and high-impact are different axes. The axis of compute and the axis of application domain move separately.

There is a reason why one cannot simply feel reassured by this distinction. The risks of high-performance models do not remain solely inside the model. The pattern of risk changes depending on the context in which the model is deployed. If user business operators using high-performance models via APIs cannot look into the internal operating principles of the model, limitations arise in identifying and managing risks. In cases where the developer business operator is overseas and a domestic representative is designated merely as a formality, questions remain open as to whether effective management and supervision are possible.

7. The Difficulty of Counting FLOPs

For the criterion of 10 to the 26th power to function, it must be possible to measure a business operator's cumulative training compute. Measurement is not easy.

The non-disclosure of training compute is a hurdle. None among OpenAI, Google, Anthropic, or Meta has officially announced the exact training FLOPs of their models. Independent research organizations

like Epoch AI back-calculate by combining published papers, interviews, and industry testimony. It is a method of calculating by weaving together model parameter count, training token count, GPU types and quantities, and training duration, but if the premises change, the results differ. There are cases where different researchers produce different estimates for the same model.

The scope of training is also ambiguous. Should we count only the compute used in pre-training, or should post-training and Reinforcement Learning from Human Feedback (RLHF) also be included? Should compute used for fine-tuning be added? The Enforcement Decree states "cumulative compute used for training" and delegated the specific calculation method to a public notice. The draft public notice issued for advance administrative notice appears to contain a direction reflecting a broad range of training-related computations, but where to draw the boundary in practice between pre-training and post-training remains in the process of being refined.

The unit of measurement itself is also not clear. FLOPs represent theoretical numbers of operations. Actual operational efficiency on hardware varies depending on GPU type, parallel processing method, and precision (whether floating-point 16-bit or 32-bit) settings. There

is also no consensus on how to handle the gap between theoretical FLOPs and measured FLOPs.

Even if a business operator voluntarily reports the compute of its model, it is difficult to externally verify that figure. It is practically not easy for regulatory authorities to directly verify GPU logs, power consumption records, and data center operation records during the training process. A structure is created where regulatory authorities rely on business operators' self-reporting.

This difficulty is not a challenge unique to Korea. Under the EU AI Act, the same discussion continues over calculating the training compute of GPAI models. The GPAI Code of Practice requires including training compute in technical documentation, but failed to unify the calculation method. In addition to the numerical criterion of 10 to the 25th power FLOPs, the EU sought to supplement the limitations of the numerical criterion by concurrently specifying a qualitative criterion of "equivalent capability."

The fact that Korea's Enforcement Decree combined qualitative requirements—the application of cutting-edge technology and concerns of significant risk—with the compute requirement can also be read as preparation for the limitations of compute as a single

indicator. However, what public notices can alter is limited to the calculation method of compute. Introducing new indicators such as parameter count or user count as target criteria requires undergoing an amendment to the Enforcement Decree.

8. An Empty Box

Very few business operators currently stand above the number 10 to the 26th power. Domestically, there are no corresponding companies. Among the latest models of global Big Tech, there are models evaluated by independent estimates as having already crossed this line, but there are no figures officially confirmed by developers. This criterion was set with the intention of giving breathing room to the Korean AI industry, and as far as domestic companies are concerned, that intention has been achieved.

However, compute increases every year. According to Epoch AI's tracking, the training compute of frontier models has doubled approximately every six months. Following this trend, 10 to the 26th power will before long become the standard for frontier models. Today's exception becomes tomorrow's norm.

Amending the Enforcement Decree can lower the threshold. The compute calculation method can be modified by public notice.

However, changing the target criteria itself is a matter of amending the Enforcement Decree, not a public notice. Amending an Enforcement Decree usually requires 40 days or more of advance notice of legislation, regulatory review, and State Council approval. In a market where the training compute of frontier models jumped several-fold in just two years after GPT-4 emerged, whether regulatory authorities can keep pace with that speed remains to be seen.

The category of high-performance artificial intelligence is still an empty box for domestic companies. The moment this box is filled will not be the day a Korean AI company's model reaches 10 to the 26th power. It will be the day global Big Tech models—which have crossed the compute threshold and satisfied the remaining two requirements—land in the Korean market. On that day, regulated entities enter the box that had been empty.

Chapter 6. Step 4 Assessment: Are We an Overseas Business Operator Required to Designate a Domestic Agent?

It is the spring of 2026, in a conference room of a shared office in Gangnam, Seoul. The Vice President of Asia-Pacific Legal Affairs for an AI company who flew in from the US headquarters sits across from a partner attorney at a Korean law firm. On the screen, Article 36 of the AI Basic Act is displayed: "An AI business operator without an address or place of business in Korea whose number of users, revenue, etc. meet the standards prescribed by Presidential Decree shall designate a domestic agent in writing and report such designation to the Minister of Science and ICT." The Vice President asks: We opened a branch in Korea; is this sufficient? The attorney shakes his head. Having a branch office and designating a domestic agent required by law are separate matters.

1. What the Legal Provision States

Article 36 consists of three paragraphs.

Paragraph 1 sets forth the conditions for the obligation and the role of the agent. If an AI business operator without an address or place of business in Korea meets the standards set by the Enforcement Decree, it must designate in writing a person to act on its behalf for

the matters enumerated by law and report this to the Minister of Science and ICT.

The domestic agent acts on behalf of the business operator in three matters: submitting the results of fulfilling the obligation to ensure safety for high-performance AI (Article 32, Paragraph 2); requesting confirmation of whether an AI system constitutes a high-impact AI (Article 33, Paragraph 1); and supporting the implementation of measures to ensure safety and reliability concerning high-impact AI (each subparagraph of Article 34, Paragraph 1). The third role includes inspecting the currency and accuracy of safety assurance documents.

Paragraph 2 defines the qualifications. A domestic agent must be a person who has an address or place of business in Korea.

Paragraph 3 is an attribution provision. If a domestic agent violates this Act in connection with the three roles above, the party deemed to have committed that act is not the agent, but the overseas business operator that designated the agent. The agent's violation is attributed to the principal. Appointing an agent does not transfer legal liability to the agent.

2. Four Criteria: Meeting Just One Triggers the Obligation

Article 29 of the Enforcement Decree establishes the criteria.

Subparagraph 1. An AI business operator whose total revenue for the preceding year is 1 trillion KRW or more. For a corporation, this is based on the preceding business year. This revenue is not limited to revenue from the AI business alone, but represents the total revenue of the relevant business operator. As an overseas business operator, this effectively means global revenue. Foreign currency revenue is converted into KRW at the average exchange rate of the preceding year (Article 29, Paragraph 2 of the Enforcement Decree).

Subparagraph 2. An AI business operator whose preceding year's revenue in the AI service sector is 10 billion KRW or more. While Subparagraph 1 examines total company revenue, Subparagraph 2 looks strictly at the AI service sector.

Subparagraph 3. An AI business operator whose average daily domestic users for its AI products and services during the preceding 3 months, as of the end of the preceding year, is 1 million or more.

Subparagraph 4. An AI business operator that has been subjected to an administrative fine pursuant to Article 43, Paragraph 1,

Subparagraph 3. Article 43, Paragraph 1, Subparagraph 3 covers cases where an administrative fine was imposed for failing to comply with a suspension order or corrective order. This provision means that once an overseas business operator receives administrative sanctions, it is obligated to designate a domestic agent thereafter, regardless of its subsequent revenue or user count.

Satisfying even one of the four criteria triggers the obligation. The phrase "falling under any of the following subparagraphs" forms the legal basis for this.

3. Do Global Big Tech Companies Qualify?

Reports indicated that OpenAI's annual recurring revenue (ARR) exceeded \$3.4 billion as of mid-2024. Although ARR is not identical to the concept of preceding year's revenue referred to in the law, since it exceeds 4 trillion KRW, it is reasonable to consider that it satisfies the Subparagraph 1 criterion (total revenue of 1 trillion KRW). The monthly visit volume of ChatGPT in Korea is known to be in the tens of millions. It is also highly likely to meet the Subparagraph 3 criterion (1 million average daily domestic users).

Google's parent Alphabet exceeded \$350 billion in total revenue in 2024. It satisfies the Subparagraph 1 criterion. As it provides Gemini

and Google Cloud AI services in Korea, there is also a possibility of satisfying the Subparagraph 3 criterion. Google has a Korean legal entity named Google Korea LLC. It might not meet the prerequisite of being an "AI business operator without an address or place of business in Korea." This is where the issue becomes complicated.

Meta's total revenue for 2024 was \$164.5 billion. It satisfies the Subparagraph 1 criterion. It offers Meta AI services in Korean. Meta maintains a corporate entity in Korea named Facebook Korea.

Anthropic's 2024 revenue is estimated to be around \$1 billion. When converted using the preceding year's average exchange rate, it is highly likely to exceed 1 trillion KRW, which could fall under the Subparagraph 1 criterion. As the number of Claude users in Korea has not been disclosed, it is difficult to determine whether it meets the Subparagraph 3 criterion. Anthropic opened a Seoul office in June 2026 and appointed a representative director for Korea. If this Seoul office constitutes a place of business as referred to in Article 36, Anthropic would fall outside the prerequisite of being an "AI business operator without an address or place of business in Korea." This is a matter where the conclusion may vary depending on which corporate entity is the contracting party for the AI service.

Latecomers such as xAI and Mistral AI may not reach the 1 trillion KRW total revenue threshold (Subparagraph 1). However, the threshold for the 10 billion KRW AI service sector revenue (Subparagraph 2) is significantly lower, meaning they may already qualify under it. One must not forget that satisfying even one of the four criteria gives rise to the obligation.

4. Is a Domestic Agent Unnecessary if There Is a Korean Branch?

Article 36, Paragraph 1 applies to business operators "without an address or place of business in Korea." If an entity has a corporate body or a place of business in Korea, it falls outside the scope of application of this provision altogether. There is no separate obligation to designate a domestic agent.

However, there are points to caution.

Even if a Korean branch exists, the key issue is whether that branch is the entity actually providing the AI service. Although Google Korea is a legal entity in Korea, the answer may differ depending on whether the contracting party providing Google's AI services to Korean users is Google Korea or the US headquarters, Google LLC. If the contracting party under the Terms of Service is the US entity and the Korean entity performs only marketing or sales support, the AI business

operator providing the AI service is the US entity. The next questions are whether that US entity can be considered to lack an address or place of business in Korea, and whether the address of the Korean subsidiary constitutes a place of business of the parent company.

This interpretation has not yet been settled. The standards will become more concrete as authoritative interpretations from the Ministry of Science and ICT or consultation cases from the AI Basic Act Support Desk accumulate. In practice, even if a Korean legal entity exists, it is a safe choice to examine the contracting party and providing entity of the AI service and, if uncertain, designate a domestic agent.

5. Comparison with the Authorized Representative of the European Union

Article 22 of the EU AI Act requires providers of high-risk AI systems established outside the European Union to designate an authorized representative established in the EU by written mandate prior to placing the system on the EU market. Article 54 separately imposes a similar obligation to designate an authorized representative on extraterritorial providers of general-purpose AI models.

Placing the Korean and EU systems side-by-side reveals a similar structure: requiring offshore business operators to establish an

onshore point of contact, and appointing an agent via written mandate. However, Korea has included an attribution provision (Article 36, Paragraph 3) that treats violations by the agent as acts of the principal, along with a prior reporting procedure, whereas the EU utilizes a structure where the mandate is presented upon request by competent authorities. Scope and depth also differ.

An authorized representative in the EU must keep a copy of the EU declaration of conformity and technical documentation, provide them to competent authorities upon request, respond to requests from market surveillance authorities, and cooperate with corrective actions. When regulatory authorities cannot reach an offshore business operator, the authorized representative effectively stands in for that operator within the EU regulatory framework.

The role of a domestic agent in Korea is limited to three matters: submitting the results of fulfilling safety assurance obligations for high-performance AI, requesting confirmation of high-impact AI, and supporting the implementation of measures to ensure safety and reliability for high-impact AI. General obligations such as transparency obligations (Article 31) or prior notification do not fall within the statutory scope of representation for domestic agents.

Application criteria also differ. In the EU, an offshore business operator placing a high-risk AI system or general-purpose AI model on the EU market incurs the obligation to designate an authorized representative regardless of revenue or user scale. Korea has established quantitative criteria: total revenue of 1 trillion KRW or more, AI revenue of 10 billion KRW or more, and average daily users of 1 million or more. Korean law exempts small overseas business operators from this obligation and imposes it only on large-scale business operators. This design realizes a form of the proportionality principle through quantitative thresholds.

The difference in sanctions upon violation is dramatic. Under the EU AI Act, the maximum administrative fine for failing to designate an authorized representative is up to 15 million EUR or 3 percent of total worldwide annual turnover, whichever is higher. Under South Korea's AI Basic Act, the upper limit of the administrative fine for failing to designate a domestic agent is 30 million KRW. To a global big tech company, 30 million KRW is not a figure that affects business decision-making.

6. China: Regulating Without Extraterritorial Application Provisions

China does not have a general declaration of extraterritorial application for artificial intelligence.

Article 2 of the Interim Measures for the Management of Generative Artificial Intelligence Services (enforced in August 2023) merely states that it "applies to the provision of generative artificial intelligence services to the public within the territory of the People's Republic of China." There is no declaration such as Article 4, Paragraph 1 of Korea's AI Basic Act stating that "this Act applies even to acts performed outside the country if they affect the domestic market or users."

The absence of a general declaration does not mean there are no provisions targeting foreign-based services. Article 20 of the same Measures stipulates that if generative AI services provided from abroad into the territory of China violate laws and regulations, the national cyberspace department shall notify relevant bodies to take technical measures and other necessary actions. China resolves jurisdictional issues through this provision combined with technical blocking. The Great Firewall physically blocks access to unregistered foreign AI services. To offer services within China, one must complete algorithm filing (备案) and pass security assessments. If a service is not registered, access itself is blocked. Instead of declaring extraterritorial application through statutory text, it defines legal boundaries using technical infrastructure.

There is also no system corresponding to Korea's domestic agent designation obligation in China. If an overseas business operator wishes to provide AI services in China, established practice virtually requires establishing a Chinese legal entity or entering into a joint venture with a local Chinese partner. It is not a structure where appointing a single agent solves the issue.

Korea declares extraterritorial application through legal provisions and requires a domestic agent as an onshore contact point. The EU requires authorized representatives across a broader scope, backed by higher administrative fines. China does not include a general extraterritorial application declaration, but instead erects physical barriers through technical blocking and prior filing. Three jurisdictions have deployed different tools to address the same problem.

7. The Weight of a 30 Million KRW Administrative Fine

Failure to designate a domestic agent results in an administrative fine of up to 30 million KRW pursuant to Article 43, Paragraph 1, Subparagraph 2.

Let us examine the implications of this amount. To global big tech companies generating tens of trillions of KRW in annual revenue, 30

million KRW is closer to a fee than a sanction. Compared to the EU AI Act, which can impose administrative fines of up to 15 million EUR (approx. 23 billion KRW) for similar violations, Korea's administrative fine has weak economic deterrent power against violations.

That being said, it is difficult to conclude that 30 million KRW is entirely meaningless.

Let us look back at Article 29, Paragraph 1, Subparagraph 4 of the Enforcement Decree. An AI business operator that has been subjected to an administrative fine pursuant to Article 43, Paragraph 1, Subparagraph 3 becomes subject to the domestic agent designation obligation thereafter, regardless of subsequent revenue or user numbers. Once an administrative fine is imposed for failure to comply with a suspension order or corrective order, the domestic agent obligation is not lifted. Rather than the 30 million KRW figure itself, attention should be paid to the structure where accumulating administrative sanction records expands legal obligations.

Even if the Ministry of Science and ICT attempts to conduct a fact-finding investigation (Article 40) while a domestic agent remains undesignated, there is no physical contact point through which to reach the target overseas business operator. The obligation to

designate a domestic agent is a mechanism designed not around the size of the administrative fine, but around securing a minimum channel through which the Korean government can communicate with overseas business operators.

A grace period applies. The imposition of administrative fines under the AI Basic Act is suspended for at least one year, until after January 2027. Unless there is serious infringement of human rights or loss of life, administrative guidance and corrective recommendations take priority. For overseas business operators, this grace period is the time to designate a domestic agent and organize their reporting system.

8. Practical Assessment Workflow

The following summarizes the step-by-step process for an overseas AI business operator to evaluate its obligations.

Does it have an address or place of business in Korea? If so, it is not subject to Article 36. There is no need to separately designate a domestic agent. However, if the contracting party for the AI service is an overseas entity and the Korean entity functions only in a supporting capacity, the overseas entity may still be subject to application.

If subject to application, does it meet any of the four criteria under Article 29 of the Enforcement Decree? Total revenue for the preceding year of 1 trillion KRW or more. Revenue in the AI service sector of 10 billion KRW or more. Average daily domestic users during the preceding 3 months of 1 million or more. History of being subjected to an administrative fine for failure to comply with a previous suspension order or corrective order. Meeting even one criterion requires designating a domestic agent.

If the criteria are met, a domestic agent must be designated in writing, and the designation reported to the Minister of Science and ICT. The domestic agent must be a person who has an address or place of business in Korea. Law firms, consulting firms, or domestic affiliates may serve as agents.

Obligations continue after designation. The domestic agent acts on behalf of the business operator in submitting the results of fulfilling safety obligations for high-performance AI, requesting confirmation of high-impact AI, and supporting the implementation of measures to ensure safety for high-impact AI. If the domestic agent violates the law within this scope, liability is attributed to the overseas headquarters.

Let us return to that Vice President in the Gangnam conference room. If a Korean branch exists, the headquarters might not be subject to Article 36 itself. However, if the contracting party under the AI service Terms of Service is the US headquarters, the headquarters may still be a business operator "without an address or place of business in Korea" despite having a Korean branch. The statutory text is short. The scope of interpretation is broad. In the EU, failure to designate an authorized representative can incur fines of tens of billions of KRW, whereas in Korea, it is 30 million KRW. A large gap in monetary amounts does not mean the obligation is light. The domestic agent is the channel through which the Korean government can reach overseas business operators; if that channel is blocked, neither fact-finding investigations nor corrective orders can be served. The essence of this provision lies not in the magnitude of the administrative fine, but in whether the minimal infrastructure required to continue operating in the Korean market has been established.

Part 3 Basic Obligations of All Business Operators: Ensuring Transparency

Chapter 7. Prior Notification Obligation: When, What, and How to Notify

One day in January 2026, in a meeting room of a fintech company in Gangnam, Seoul. The legal team leader projected a screenshot of the company's loan comparison app onto the screen. The app analyzes users' credit information and recommends loan products in ascending order of interest rates. The recommendation engine is a machine learning model. The legal team leader asks: "According to Article 31, Paragraph 1 of the AI Basic Act, we must notify users in advance that AI is being used. But when exactly is 'in advance'? Is it when downloading the app for the first time? Is it upon user registration? Or is it right before displaying the recommendation results?" The development team leader continues: "Is the notification method a pop-up, the terms of service, or fine print at the bottom of the screen? And is it enough to simply write 'AI is being used,' or do we also need to specify what kind of AI it is and what data it was trained on?" The legal team leader read through the legal provisions but could not find an answer. This was not a problem unique to this company. Of the online inquiries received by the AI Basic Act Support Desk, 56.4 percent were related to transparency obligations.

1. What the Law Requires

Article 31, Paragraph 1 of the AI Basic Act is a single sentence:

"Where an AI business operator intends to provide a product or service using high-impact artificial intelligence or generative artificial intelligence, it shall notify users in advance of the fact that the product or service is operated based on the relevant artificial intelligence."

Deconstructing this sentence reveals four elements.

The subject of the obligation is the AI business operator. This applies to both development business operators and utilization business operators. It covers not only companies that directly develop and serve AI models, but also companies that adopt third-party AI to provide products or services under their own brand.

The target AI consists of two types: high-impact artificial intelligence and generative artificial intelligence. On one side is high-impact AI used in the ten domains specified by law—such as energy supply, healthcare, recruitment, credit assessment, traffic management, public services, and student evaluation—as well as domains prescribed by Presidential Decree. On the other side is generative AI that creates text, images, audio, and video. Business

operators using AI that does not fall under these two categories are not subject to the prior notification obligation under this Article. Machine learning models that classify customer inquiries or logistics optimization algorithms do not fall under high-impact domains nor do they generate content, so they are excluded from the scope of prior notification under Article 31, Paragraph 1.

The content of the notification is clear: "the fact that the product or service is operated based on the relevant artificial intelligence." It is the mere fact that AI is being used. What model is used, what data it was trained on, or what percentage of accuracy it achieves falls outside the scope required by this Article.

The timing is "in advance." Users must be notified before using the product or service.

In case of violation, an administrative fine of up to 30 million KRW is imposed pursuant to Article 43, Paragraph 1, Item 1. Under the AI Basic Act, only three obligations are subject to administrative fines: failure to perform prior notification, failure to designate a domestic agent, and failure to comply with suspension or corrective orders. Prior notification is one of those three. While there is debate over whether 30 million KRW acts as a sufficient deterrent for large

enterprises, it is clear that it was included as a subject of administrative fines from the very first year of the law's implementation. However, a grace period of at least one year applies, so administrative guidance and corrective recommendations take precedence unless it involves serious matters such as loss of human life or large-scale human rights violations.

2. When Is "In Advance"?

Article 31, Paragraph 1 of the Act simply states "in advance." Article 23 of the Enforcement Decree lists notification methods but does not define the specific timing of "in advance." This omission appears to be intentional. Because the forms of products and services employing AI are so diverse, fixing a single point in time would result in cases that do not fit reality.

In legal practice, "in advance" is interpreted in three ways depending on the type of service.

For services requiring a contract, prior to contract execution is the natural timing. In the case of a loan comparison app, a notice stating "This service recommends products based on artificial intelligence" must be provided before the user agrees to the terms of service. Belatedly notifying users after contract execution deprives them of

the opportunity to choose an AI-based service, failing to align with the intent of "in advance."

For services accessed without a separate contract, such as websites or apps, the baseline is the time of initial use. It is sufficient if users can confirm the fact of AI-based operation when they first open the service screen. Discovering it deep within the settings menu after using the service for a while does not constitute "in advance."

For offline products, it means prior to purchase or prior to initial use. A notice must be present on the packaging or in-store display of home appliances equipped with AI.

The Transparency Assurance Guidelines support this interpretation. Although the guidelines are not legally binding, they serve as practical guidance formulated by the Ministry of Science and ICT to flesh out the intent of Article 31, making them the primary document business operators should reference when designing compliance. The guidelines present the principle that notification must occur "before the user decides to use the product or service." The benchmark of assessment is the user's "decision." The opportunity for users to choose to use the service knowing that it is AI-based must be ensured.

3. Four Means of Notification

Article 23, Paragraph 1 of the Enforcement Decree enumerates four notification methods.

The first method is stating directly on the product or stating it in contracts, user manuals, terms of service, etc. Writing "This software is operated based on artificial intelligence" in a contract when supplying B2B software belongs to this type. Printing "Artificial intelligence-based voice recognition product" on the box of an AI speaker is similar. The advantage is that evidence remains and notification is naturally completed during the contract execution process. The disadvantage is that if buried among long sentences in the terms of service, users may fail to actually perceive it.

The second method is displaying on the user's screen or terminal. Pop-ups, banners, in-app notifications, and notices at the top of web pages fall into this category. An example is having a loan comparison app display a pop-up stating "This app provides an artificial intelligence-based recommendation service" when launched for the first time and prompting the user to tap a confirmation button. It is a foolproof method in that it catches the user's eyes directly. However, it has limitations, as accumulated pop-up fatigue may lead users to adopt the habit of closing them without reading.

The third method is posting in an easily recognizable manner at the place where the product or service is provided. A proviso in parentheses states, "including places within a scope reasonably related to the place concerned." Posting a sign at the entrance of an unmanned convenience store stating "This store operates AI-based customer movement analysis" is one example. A company utilizing an AI interview system posting "Artificial intelligence analysis is included in the interview process" on its job recruitment page also fits this type.

Lastly, there is the method recognized by the Minister of Science and ICT. This is an open-ended provision. As technology advances, notification methods may also change, serving as a mechanism to accommodate new methods that are difficult to fit into the preceding three.

The Enforcement Decree leaves the choice among these four methods to the business operator. Operators may use just one or combine several. However, regardless of the method chosen, users must be able to actually perceive it. Writing it in 8-point font at the end of page 47 of the terms of service, even if formally constituting "stating," fails to fulfill the intent of "prior notification" meant by the

guidelines.

4. To What Level Must Notification Be Provided?

The content of notification required by Article 31, Paragraph 1 of the Act is "the fact that it is operated based on the relevant artificial intelligence." The mere fact that AI is used. This Article asks for nothing beyond that.

Specifically, information such as the following is not subject to prior notification under Article 31, Paragraph 1: which model is used (whether GPT-4 or Gemini), what data it was trained on (whether public data or internally collected data), what the model's accuracy or error rate is, or what the operational principles of the algorithm are.

"We use AI" is the baseline floor for notification.

However, another obligation exists above the baseline floor. For business operators operating high-impact artificial intelligence, Article 34, Paragraph 1, Item 2 imposes a separate obligation of explanation. They must establish and implement explanation measures regarding the final results derived by artificial intelligence, the primary criteria used in deriving final results, an overview of the

training data used in the development and utilization of artificial intelligence, etc., "to the extent technically feasible." Article 27, Paragraph 1 of the Enforcement Decree stipulates that the main details of these explanation measures must be posted on the business operator's internet website or other media. Matters constituting trade secrets may be excluded.

In summary: prior notification under Article 31, Paragraph 1 is the notification of the fact that "AI is being used." The explanation obligation under Article 34, Paragraph 1, Item 2 is an explanation of "how AI makes decisions." The former applies to all high-impact AI and generative AI, whereas the latter applies only to high-impact AI. The former must be fulfilled prior to use, while the latter must be continuously fulfilled during service operation.

The basic principle provision of Article 3, Paragraph 2 goes one layer further: "Affected persons shall be able to receive clear and meaningful explanations regarding the major criteria and principles utilized in deriving the final results of artificial intelligence, to the extent technically and reasonably feasible." It is a basic principle that affected persons should be able to receive explanations. Although it is not a right attached to separate claim procedures or sanctions, it is

a provision clarifying the direction that when high-impact AI makes a decision unfavorable to an individual, that individual should be able to inquire why such a decision was made.

Prior notification, explanation obligations, and basic principles of explanation provision. Three layers of transparency mechanisms overlap, and prior notification is the first layer among them. This is because if one does not know the fact that "AI is being used," one can neither request nor receive an explanation.

The Transparency Assurance Guidelines advise business operators to draft prior notification text as clearly and easily understandable as possible. Rather than the phrase "AI-based," it is desirable to include the specific functionality of the service, such as "Artificial intelligence (AI) provides recommendation results in this service." There are also guidelines to avoid technical jargon and consider users' age and digital literacy. Article 23, Paragraph 3 of the Enforcement Decree stipulates that in deepfake-related notifications, the notification method must be designed "taking into account the age, physical, and social conditions of main users," a principle that is also worth referencing for general prior notification.

5. Output Labeling and Deepfake Labeling

There are two more obligations in the same legal article that must be distinguished from prior notification under Article 31, Paragraph 1.

Article 31, Paragraph 2 is the output labeling obligation. A business operator providing generative artificial intelligence or products or services using it must indicate the fact that its output was generated by generative artificial intelligence. If prior notification is a notice about the service itself, output labeling is a label on each individual output item. Attaching an "AI Generated" label to a photo created by an AI image generation service is an example of output labeling.

Article 23, Paragraph 2 of the Enforcement Decree divides labeling methods into two categories: methods perceptible to humans (visible labels, text phrases, audio guidance) and methods readable by machines (metadata, digital watermarks). When using machine-readable methods, the fact that the output was generated by generative AI must be provided at least once through notice text or audio guidance. A watermark readable only by machines is not sufficient on its own.

Article 31, Paragraph 3 goes one step further in strictness. Where synthetic audio, images, or video that are difficult to distinguish from reality are created using an AI system, notice or labeling must be

provided in a manner that users can "clearly perceive." This provision targets deepfakes. However, in cases constituting artistic or creative expressions, notification may be provided in a manner that does not interfere with exhibition or appreciation. Since overlaying an "AI Generated" watermark across the entire screen during a scene of an AI movie would spoil the artwork, it means it may be substituted by displaying it in the end credits or work description.

Article 23, Paragraph 3 of the Enforcement Decree adds two criteria to deepfake labeling: display in a manner where contents can be easily verified through sight, hearing, or software; and take into account the age and physical and social conditions of main users. For services primarily used by the elderly, this implies increasing font size or providing audio guidance.

6. Exceptions to Notification Requirements

Article 23, Paragraph 4 of the Enforcement Decree establishes three exceptions.

The first is where the fact that high-impact artificial intelligence or generative artificial intelligence is used is obvious in light of the name of the product or service, user screen, product exterior, or text displayed on outputs. Services whose names themselves reveal that

they are AI, such as "ChatGPT," "DALL-E," or "AI Translator," do not require separate notification. Because users access the service already knowing it is an AI-based service, the purpose of notification has been achieved.

The second is where it is used solely for the internal business purposes of the AI business operator. If employees within a company use AI for internal analysis or decision-making support, but the results are not directly provided to external users, there is no obligation for prior notification. However, this exception must be interpreted narrowly. If decisions are made regarding users based on AI analysis results (for instance, screening resumes with AI and subsequently sending pass or fail notifications to applicants), that AI is not "used solely for internal business purposes." This is because the results impact users.

Lastly, cases where the Minister of Science and ICT determines by public notice that an exception to application is necessary, taking into consideration the type and characteristics of the product or service, content of outputs, form of use, level of technology, etc. As of July 2026, this public notice has not yet been issued.

There is a point of caution when reading exception provisions. These exceptions are not limited to prior notification (Article 31, Paragraph 1), but may apply in whole or in part to output labeling (Paragraph 2) and deepfake labeling (Paragraph 3). However, falling under an exception does not mean exemption from transparency obligations as a whole. The explanation obligation (Article 34) and risk management obligations for high-impact AI business operators remain separate and intact.

7. European Union: Transparency for All AI Systems

Article 50 of the EU AI Act demands transparency over a broader scope than South Korean law.

Article 50(1). Where an AI system is designed to interact directly with natural persons, providers shall inform those natural persons that they are interacting with an AI system. "Cases where it is obvious from the circumstances and context of use" are exceptions. This means that when a chatbot converses with a customer, it must disclose "I am an AI." While sharing the same principle as South Korean law, the scope of application differs. South Korea targets only high-impact AI and generative AI for notification. The EU targets all AI systems that interact with natural persons, regardless of risk classification.

Article 50(2). Providers of AI systems that generate synthetic audio, image, video, or text content must mark the outputs in a machine-readable format to indicate AI generation. South Korean law recognizes machine-readable methods, but considers them insufficient on their own, requiring that human-perceptible guidance be provided at least once alongside them. The EU places machine-readable marking as the provider's basic obligation, while leaving human-perceptible disclosure to the deployer.

Article 50(3). Deployers of emotion recognition systems or biometric categorization systems must inform natural persons exposed to such systems of their operation. South Korean law has no separate notification obligation for emotion recognition AI.

Article 50(4). Deployers of AI systems that generate or manipulate image, audio, or video content constituting deepfakes must disclose that the content has been artificially generated or manipulated. Disclosure obligations also apply where text published to inform the public on matters of public interest is generated or modified using AI. However, where AI-generated content undergoes human review and editorial responsibility is assumed (such as a news reporter editing an AI draft before publishing an article), the text disclosure obligation is

exempted.

South Korean law does not have this "editorial responsibility" exemption. In principle, outputs created by generative AI must be labeled.

South Korea looks at impact. The EU looks at interaction. South Korean law divides transparency obligations based on the type of AI (high-impact, generative). EU law divides obligations based on how the AI system contacts humans. As a result, the EU's net is cast wider. Low-risk recommendation algorithms that are not subject to prior notification in South Korea become subject to notification in the EU if they directly interact with users.

The gap in administrative fines remains stark. South Korea's administrative fine for violating prior notification is up to 30 million KRW. According to Article 99 of the EU AI Act, fines for violating Article 50 reach up to 15 million EUR or 3 percent of total worldwide annual turnover, whichever is higher. That is a difference of more than 750 times.

Transparency obligations under the EU AI Act (Article 50) apply from August 2, 2026. Since transparency obligations under South Korea's

AI Basic Act (Article 31) entered into force on January 22, 2026, the EU side is about six months later. Both laws take effect around the same period. Although the EU decided to postpone high-risk AI obligations until December 2027 through the Digital Omnibus amendment, it left the timeline for transparency provisions—such as chatbot disclosure and deepfake labeling—unaltered. Transparency was deemed an obligation that could not be deferred.

8. China: Register, Label, and Control

China's AI transparency obligations are designed differently from those of South Korea or the EU. Transparency provisions are scattered across sectoral regulations rather than contained in a single comprehensive statute.

The Administrative Provisions on Algorithm Recommendations for Internet Information Services (enforced March 2022) impose user notification obligations on algorithm recommendation service providers. Providers must notify users of the fact that algorithm recommendation services are being provided and guarantee users the option to turn off personalized recommendations. Up to this point, it is similar to prior notification under South Korean law. The point of divergence is filing/registration (备案). Providers of services with public opinion attributes or social mobilization capabilities must

file details such as the type of algorithm, application fields, and self-assessment results with the Cyberspace Administration of China (CAC). Filing must take place within 10 business days from the date service provision begins, and any changes to filed content require updated filing within 10 business days. South Korea has no obligation to register AI with the government.

The Interim Measures for the Management of Generative Artificial Intelligence Services (enforced August 2023) go one step further. To provide generative AI services with "public opinion attributes or social mobilization capabilities" to the public, providers must apply for a security assessment and complete algorithm filing prior to launch. Service providers must notify users of the fact that generated content was created by AI and comply with "core socialist values." South Korean law contains no provisions controlling the substance of content itself.

The Measures for Labeling AI-Generated Synthetic Content, which took effect on September 1, 2025, precisely standardized labeling obligations. Service providers must embed machine-readable implicit labels (metadata, digital watermarks) into AI-created content, and for content types capable of causing confusion like

deepfakes, they must also provide human-visible explicit labels (visible labels, text phrases). The mandatory national technical standard GB 45438-2025 standardized the format and specifications of labeling. South Korean law leaves labeling methods to the business operator's discretion without setting specific specifications. China goes so far as to have the state set labeling specifications.

Filing statistics illustrate the scale of this regime. As of August 2024, over 190 large language models had completed filing, and by November 2025, 611 generative AI services were registered. The total number of algorithm filings is reported to have exceeded 5,000 cases.

Regulations continue to expand even now. Jointly promulgated by five authorities including the CAC in April 2026 and enforced from July 15 of that year, the Administrative Provisions on Personified AI Interaction Services established an identity disclosure obligation requiring AI services that converse and exchange emotions like humans to clearly inform users that their counterpart is artificial intelligence, not a human being. Moving a step beyond labeling content, the intent is to prevent AI from posing as a human and swaying user emotions and judgment.

Looking across the three jurisdictions, the meaning of transparency differs. South Korea notifies users. The EU notifies users and labels outputs. China notifies users, labels outputs, and files with the government. South Korea's transparency points toward the user. China's transparency points toward both the user and the state.

Differences also exist in enforcement intensity. South Korea's maximum administrative fine of 30 million KRW virtually boils down to administrative guidance during the grace period. China, on the other hand, is reported to have deleted hundreds of thousands of illegal content items and penalized numerous AI accounts and services during the internet authority's campaign against AI technology abuse in 2025. De facto severe sanctions in the form of service suspension follow.

9. Practical Design Challenges

After reading through legal articles, the Enforcement Decree, and guidelines, design challenges remain before business operators.

How should notification text be written? A single line stating "This service is operated based on artificial intelligence" is the minimum. However, whether this single line conveys meaningful information to users is another matter. For a loan product recommendation app,

writing "Artificial intelligence analyzes your credit information to recommend loan products in this app" aids user understanding. The guidelines also recommend notifications linked to specific service functionality.

Where should it be placed? Legally, putting a single line in the terms of service suffices. Practically, placing it where the user's line of sight lands on the service screen prevents disputes. This is because "whether the user actually perceived it" is far more likely to become the issue in a dispute than the mere fact that "it was notified." Which location is appropriate—the consent screen during registration, the top banner of the main service screen, or supplementary explanations on the result display screen—depends on the nature of the service.

Is notifying once sufficient? The law merely specifies "in advance" without setting a frequency. Once notified, there is no statutory obligation to repeat it thereafter. However, if the service is updated and the role of AI changes, or if AI is newly introduced into a feature where AI was previously not used, re-notification is required. An updated notification stating "This feature has been changed to be AI-based" is necessary.

What if an operator is uncertain whether its AI qualifies as high-impact AI? Requesting confirmation from the Ministry of Science and ICT yields a response within 30 days. A one-time extension of 30 days is possible. If there is an objection to the confirmation result, a re-confirmation request can be submitted within 10 days, and the re-confirmation decision is issued within 30 days following consultation with an expert committee. Even while confirmation procedures are underway, it is safer to provide prior notification if the service is already being offered. Even if the result reveals that it is not high-impact AI, having provided notification causes no harm.

There is no single correct answer to all these questions. Article 31 and Article 23 of the Enforcement Decree merely set the framework; specific execution within that framework remains the responsibility of business operators. As of March 31, 2026, the 552 consultations received by the AI Basic Act Support Desk represent the footsteps of business operators navigating their way within this framework.

Earlier, the legal team leader in the meeting room asked: "When exactly is 'in advance'?" The law did not specify a date. Instead, it provided a principle: users shall not be deprived of the opportunity to decide whether or not to use an AI-based service. Translating that

principle to fit one's own service remains a task for business operators.

Chapter 8. Output Labeling Obligations: Watermarking, Labeling, and Special Provisions for Deepfakes

In October 2025, 12 photos of a newly built studio apartment in Gangnam, Seoul, were uploaded to a real estate brokerage platform. A spacious living room, a night view of the Han River beyond wall-to-wall glass, and an island kitchen. The photos were crisp and the lighting was natural. Inquiries poured in. However, that studio apartment was still under structural framework construction. All 12 photos were AI-generated images. The sales company did not indicate anywhere that AI had been used. The AI Basic Act, which entered into force on January 22, 2026, was not yet in effect at that time.

After this Act entered into force, circumstances change. Article 31 of the Act requires AI business operators to label outputs of generative AI. The business operator of the image generation service that produced those photos primarily bears this obligation, while whether the sales company that created and used the promotional materials with AI constitutes an addressee under this Act depends on the case. Article 23 of the Enforcement Decree specifies the methods of and exceptions to such labeling in detail. Article 31 stacks a three-layered set of obligations—prior notification, output labeling, and deepfake

notification—within a single article.

1. Three Layers of Transparency

Unpacking the structure of Article 31:

Paragraph (1) prescribes prior notification. When providing products or services utilizing high-impact AI or generative AI, the fact that the product or service is operated based on the relevant AI must be notified to users in advance. This is a notice stating, "This service uses AI." It is a designation regarding the service itself, not a label on each individual output.

Paragraph (2) prescribes output labeling. When providing generative AI or products/services utilizing it, the fact that the output was generated by generative AI must be labeled. If Paragraph (1) is a sign placed at the entrance of a service, Paragraph (2) is a tag attached to every output produced by the service. A report generated by ChatGPT, an image drawn by Midjourney, and music composed by Suno AI each must be labeled with "Created by AI."

Paragraph (3) sets forth a special provision for deepfakes. When providing outputs such as virtual sounds, images, or videos that are difficult to distinguish from reality using an AI system, notice or

labeling must be provided in a manner that users can clearly perceive. This is a requirement heightened by one level compared to Paragraph (2). Paragraph (2) says "Label it," whereas Paragraph (3) says "Notice or label it in a manner that users can clearly perceive." An additional modifier has been attached.

The three-layered obligations differ in their scopes of application. The prior notification under Paragraph (1) encompasses both high-impact AI and generative AI. Whether it is an AI system for recruitment screening (high-impact) or an image generation service (generative), users must be notified in advance that AI is being used. The output labeling under Paragraph (2) applies exclusively to generative AI. It is limited to cases where AI creates content rather than making judgments or recommendations. The deepfake special provision under Paragraph (3) covers even non-generative AI systems. Because it states "using an artificial intelligence system," even a non-generative AI becomes subject to this provision if it creates sounds or videos difficult to distinguish from reality.

2. Visible Labeling and Machine-Readable Labeling

How business operators actually implement labeling depends on Article 23 of the Enforcement Decree.

Paragraph (1) lists four methods of prior notification: stating directly on the product or specifying in contracts, user manuals, or terms of service; displaying on the user's screen or terminal; posting in an easily recognizable manner at the location where the service is provided; and other methods recognized by the Minister of Science and ICT. For an AI chatbot service, placing a banner at the top of the chat screen stating "This service is operated by artificial intelligence" is an immediately obvious method of implementation.

Paragraph (2) divides output labeling methods into two.

One is a method perceivable by humans: visible labeling. Text stating "AI-generated," a watermark logo inserted in the corner of an image, or subtitles placed at the bottom of a video fall under this category.

The other is a method readable by machines: invisible labeling. This involves embedding information into a file's metadata stating "This content was generated by AI," or inserting a digital watermark invisible to the human eye into an image. Separate software is required to read metadata or watermarks; humans cannot verify them with the naked eye.

Here, a crucial design choice is revealed. The Enforcement Decree allows business operators to choose between the two, stating that it "may be done by any of the following methods." Performing only visible labeling is permissible. Performing only machine-readable labeling is also permissible. However, if only the machine-readable method is chosen, a condition is attached: "The fact that the output was generated by generative AI must be provided at least once through notice text, audio, or other means." Merely hiding information within metadata is insufficient; the fact of AI generation must be notified to the user at least once in a human-readable format.

Summarizing this structure: a business operator can complete compliance solely with visible labeling. Alternatively, a method combining machine-readable labeling with at least one notification statement is also possible. Relying solely on machine-readable labeling alone is not permitted.

3. Mitigation for Deepfakes Is Not an Exemption from Labeling

The deepfake special provision under Article 31, Paragraph (3) establishes a higher standard than Paragraph (2).

The phrasing "in a manner that users can clearly perceive" is key. Article 23, Paragraph (3) of the Enforcement Decree unfolds this requirement into two elements: users must be able to easily verify the content visually, audibly, or through software; and the age, physical, and social conditions of primary users must be taken into account.

The term "clearly perceive" is difficult to satisfy with machine-readable labeling alone. Even if information is embedded in metadata, ordinary users do not open metadata viewers. Even if an invisible watermark is inserted, no one can know unless detection software is executed. The danger of deepfakes lies in their speed of dissemination. Once spread, it is difficult to undo. The intent of this provision is to make it visible to human eyes in advance rather than verifying it with machines *ex post facto*.

The Act provides one mitigation. Where deepfake outputs constitute artistic or creative expressions or form part thereof, notice or labeling may be provided in a manner that does not hamper display or enjoyment. Short films or music videos created with AI fall under this case. Continuously displaying the text "AI-generated" over the main screen of a film interferes with viewing. In such cases, it means it can be substituted by specifying it in the ending credits or

posting a notice on the distribution page. However, this does not exempt the labeling itself.

4. Three Situations Where Exceptions Apply

Article 23, Paragraph (4) of the Enforcement Decree provides three exceptions. These exceptions span across all three obligations—prior notification, output labeling, and deepfake notification—and are structured so that all or part of the regulations may not apply when applicable.

Where the fact of AI utilization is obvious when comprehensively considering the name of the product or service, the user interface screen, the exterior surface of the product, text displayed on the output, etc., separate labeling may not be required. There is no need to write "AI-generated" again on an image produced by an app named "AI Image Generator." The service name itself serves as notification. This exception is a commonsense safeguard against redundancy.

Where it is used solely for the internal operational purposes of a business operator, all or part of the notice and labeling obligations may also not apply. If a marketing team generates an AI image for an idea meeting and places it on an internal slide deck, labeling is not

required. The moment that image is used for external advertising, this exception vanishes. The boundary moving from internal to external is precisely the starting line of obligation.

The third is an exception determined by notice of the Minister of Science and ICT. Additional exceptions to application may be prescribed considering the type and characteristics of products or services, the content of outputs, patterns of use, technical levels, etc. This provision is an open delegation for which specific public notices have not yet been issued.

5. What Happens in Case of Violation?

According to Article 43, Paragraph (1), Item 1 of the Act, a person who violates Article 31, Paragraph (1)—that is, the prior notification obligation—is subject to an administrative fine not exceeding 30 million KRW.

There is a notable point here. The administrative fine provision specifies only violations of Article 31, Paragraph (1). There are no provisions in Article 43 for administrative fines regarding violations of Paragraph (2) (output labeling) or Paragraph (3) (deepfake notification/labeling). Instead, Article 40, Paragraph (1) prescribes that if the Minister of Science and ICT discovers or receives a report

on a violation of Article 31, Paragraph (2) or (3), the Minister may request submission of materials or conduct a factual investigation. Article 40, Paragraph (3) provides that if a violation is recognized as a result of the investigation, a suspension order or corrective order may be issued. Furthermore, according to Article 43, Paragraph (1), Item 3, an administrative fine not exceeding 30 million KRW is imposed on anyone who fails to comply with such a suspension order or corrective order.

Tracing the structure reveals the following: failure to label output does not immediately incur an administrative fine by itself. An administrative fine arises only when the Ministry of Science and ICT investigates, orders correction, and the business operator fails to comply with that corrective order. Non-compliance with prior notification results in an immediate administrative fine. Non-compliance with output labeling passes through an intermediate step of non-compliance with a corrective order. The grace period (at least one year, deferred until after January 2027) also applies to administrative fines for prior notification. Immediate investigations are possible in major cases such as criminal misuse of deepfakes or casualties.

6. The European Union: Mandatory Machine-Readable Labeling

Article 50 of the EU AI Act occupies the same position as Article 31 of the Korean law, but takes a different approach. Its official title is "Transparency obligations for providers and deployers of certain AI systems." Composed of seven paragraphs, it is more detailed than the three layers of the Korean law.

Article 50, Paragraph 1 relates to interaction notification. Providers of AI systems must ensure that users are aware they are interacting with an AI system. An exception applies where "this is obvious from the point of view of a natural person who is reasonably well-informed." This corresponds to the prior notification under Article 31, Paragraph (1) of the Korean law, but the EU incorporated the "obvious case exception" directly into the statutory text. Korea regulates this exception in its Enforcement Decree.

Article 50, Paragraph 2 prescribes synthetic content labeling. It requires providers of AI systems to mark outputs of synthetic audio, image, video, or text content in a machine-readable format indicating that they have been artificially generated or manipulated. Technical solutions must be "effective, interoperable, robust and reliable." Recital 133 lists acceptable technologies: watermarks, metadata identification, cryptographic methods proving provenance and

authenticity, logging methods, fingerprints, etc. C2PA aligns precisely with the cryptographic methods and metadata identification mentioned in this list. This contrasts with the Korean law, which allows business operators to choose between visible labeling and machine-readable labeling. The EU established machine-readable labeling as a baseline obligation. AI systems performing only an assistive editing function or systems that do not materially alter input data are exempt from this obligation.

Article 50, Paragraph 4 sets out separate obligations for deepfakes and AI-generated text, grouping the two together. Where an AI system is used to generate or manipulate image, audio, or video content that resembles existing persons, objects, places, or events in a way that appears authentic, the fact that the content has been artificially generated or manipulated must be disclosed. This obligation is imposed on deployers. Exceptions apply when used for law enforcement purposes, cases falling under freedom of expression such as satire or parody, and artistic creations. Even for artistic expressions, labeling itself is not exempt and must be marked in "an appropriate manner that does not hamper the display or enjoyment of the work." This structure is nearly identical to the mitigation for artistic expression under Article 31, Paragraph (3) of the Korean law.

The difference lies in the explicit inclusion of satire and parody; the Korean law lacks this explicit exemption. The latter half of the same paragraph also addresses AI-generated text published for public interest information. Where human editorial review and editorial responsibility are exercised, it is exempt from disclosure obligations. This regulation keeps in mind situations where news media use AI as an assistive tool while journalists retain final editorial control.

Article 50, Paragraph 7 delegates the encouragement of codes of practice to the European AI Office. It provides for drawing up European-level codes of practice for the detection and labeling of AI-generated content and allows the European Commission to adopt implementing acts.

Summarizing structural differences between the EU and Korea:

Korea examines impact; the EU examines risk. This distinction recurs in labeling obligations as well.

Article 31, Paragraph (2) of Korean law ends with "must label." The method is chosen by the business operator in the Enforcement Decree, whether visible or machine-readable. Article 50, Paragraph 2 of the EU Act mandates machine-readable formats. Korea gives

business operators options; the EU dictates the technical implementation method.

Korea's deepfake provision (Article 31, Paragraph (3)) requires "a manner that users can clearly perceive," effectively compelling visible labeling. The EU's deepfake provision (Article 50, Paragraph 4) requires disclosure. Both sides share the recognition that machine-readable labeling alone is insufficient for deepfakes.

The transparency obligations of the EU AI Act became applicable from August 2, 2026. Although high-risk AI obligations were postponed to December 2027 due to Digital Omnibus amendments, transparency obligations (Article 50) entered into force per the original schedule. The Omnibus amendments did not alter Article 50 itself. A four-month grace period was granted for systems already placed on the market prior to August 2, 2026.

Disparities in administrative fines are confirmed in this provision as well. The EU can impose fines up to 15 million EUR or 3% of total worldwide annual turnover for violations of Article 50. Korea imposes up to 30 million KRW upon failure to comply with a corrective order.

7. China: Establishing a Two-Layered Labeling System

On September 1, 2025, the "Measures for the Labeling of Artificial Intelligence Generated and Synthesized Content" took effect in China. It was jointly promulgated by four authorities: the Cyberspace Administration of China (CAC), the Ministry of Industry and Information Technology, the Ministry of Public Security, and the National Radio and Television Administration.

Korea chooses between visible labeling and machine-readable labeling. China stacks two layers of labeling on top of each other.

It is a two-layered system of explicit labeling (显式标识) and implicit labeling (隐式标识). Explicit labeling refers to labels visible to users' eyes. Adding the text "AI-generated" on an image or displaying an AI-generated notice on screen during video playback falls under this category. Implicit labeling refers to machine-readable information embedded in a file's metadata.

The scopes of application for the two layers differ. Implicit labeling—that is, embedding AI-generation information into file metadata—is an obligation spanning synthesized and generated content as a whole. Explicit labeling is required for services outputting content that could cause confusion or misunderstanding among the public. Digital watermarking embedded directly into

content itself remains at an encouragement level rather than mandatory. Furthermore, anyone is prohibited from maliciously removing, tampering with, forging, or concealing these labels. The original text stating "shall not maliciously remove, tamper with, forge, or conceal (不得恶意删除、篡改、伪造、隐匿)" clarifies this prohibition.

In addition, the mandatory national technical standard GB 45438-2025 entered into force on the same day. This standard prescribes technical details for labeling methods. It is a mandatory standard, not a voluntary one. Apart from international standards such as C2PA or ISO, labeling of AI-generated content circulated within China must follow this national standard.

Placing the three countries side by side reveals the order of regulatory intensity.

China mandates implicit metadata labeling across content generally and adds explicit labeling to services with risks of confusion. It has even set a mandatory national standard. The EU mandates machine-readable labeling while adding disclosure obligations exclusively for deepfakes. Korea allows business operators to select either visible labeling or machine-readable

labeling. For deepfakes alone, it requires clear perceivability, in effect encouraging visible labeling.

Korea's provision of choices reflects a decision taking into account the pace of technological development and industrial burden. Mandating a specific technology method while machine-readable labeling technology is not yet mature could shift excessive costs onto business operators. Conversely, China demonstrates an intent to control even technical methods through national standards. The EU stands in between, establishing the principle of "machine-readable formats meeting state-of-the-art technology" while entrusting specific technologies to Codes of Practice.

8. Can Watermarking Technology Keep the Law's Promise?

The law requires labeling. Technology is the means to execute that labeling. As of July 2026, has watermarking technology reached the level required by law?

The Coalition for Content Provenance and Authenticity (C2PA) is a content provenance coalition jointly founded in February 2021 by Adobe, Arm, BBC, Intel, Microsoft, and Truepic. Specification version 2.1 was released in September 2024, and through successive updates, reached version 2.4 in 2026. C2PA embeds a structure called a

manifest store inside files. This structure contains information such as content generation tools, editing history, and whether AI was used to generate it, signed with X.509 certificates. Hard binding binds the entire file with a hash to detect tampering. Soft binding uses perceptual hashing or fingerprinting to track provenance even after files are re-encoded.

Adobe integrated Content Credentials into Photoshop, Lightroom, and Firefly. Microsoft applied it to Bing Image Creator and Designer. OpenAI embeds C2PA metadata into image outputs of DALL-E 3 and GPT-4o. Camera manufacturers Nikon (Z9, etc.), Sony (A9III, etc.), and Leica (M11-P) have also mounted C2PA-supported firmware. Google announced in 2024 that it would display C2PA information in search results. Work to standardize C2PA specifications into an international standard is also underway.

Google DeepMind's SynthID takes a different path. It directly embeds watermarks invisible to human eyes into the content itself. It supports images (Imagen), text (logit sampling bias in Gemini), audio, and video. It is designed to survive a certain level of compression or cropping. Meta developed Stable Signature for images and Video Seal for videos.

The limits of technology are clear.

C2PA metadata is fragile. Taking a screenshot strips away metadata. Uploading to social media platforms strips most EXIF and XMP information. While soft binding mitigates this issue, it remains vulnerable to aggressive cropping or adversarial transformations. Invisible watermarks (SynthID method) face another dilemma. Embedding a strong watermark degrades content quality; embedding a weak one allows removal by attacks such as Gaussian noise, repeated JPEG compression, or diffusion-based purification.

No single technology satisfies all conditions required by law. While C2PA metadata has a broad ecosystem, it is meaningful only if metadata is preserved throughout the entire distribution path; if major social media platforms fail to maintain it, it becomes rendered powerless. Invisible watermarking is a probabilistic detector, not a cryptographic proof.

A hybrid approach adding invisible watermarking to C2PA metadata and combining platform-level registries is currently the leading direction. However, standardization and cross-platform interoperability remain incomplete.

The fact that Article 23, Paragraph (2) of the Korean Enforcement Decree attached the condition "must be provided at least once through notice text, audio, or other means" when choosing machine-readable methods reflects the recognition that machine-readable technology alone is insufficient to protect users. Until technology can fully execute the law's promise, the analogue safeguard of human-visible labeling serves as a safety net.

9. Business Operator Practical Compliance Checklist

Synthesizing the Act and the Enforcement Decree, items that generative AI business operators must verify are as follows:

Has notification of AI-based operation been provided at the entrance of the service? Whether through terms of service or screen display, users must be able to know that AI is involved before using the service. This corresponds to Article 31, Paragraph (1) of the Act and Article 23, Paragraph (1) of the Enforcement Decree. Violation immediately subjects the operator to an administrative fine not exceeding 30 million KRW.

Have outputs generated by AI been labeled? One must select either visible labels or machine-readable watermarks. If only machine-readable labeling was selected, was notice text provided at

least once? This corresponds to Article 31, Paragraph (2) of the Act and Article 23, Paragraph (2) of the Enforcement Decree.

Are the outputs virtual sounds, images, or videos that are difficult to distinguish from reality? If so, have they been labeled in a manner that users can clearly perceive? Machine-readable labeling alone is unlikely to satisfy this standard. This corresponds to Article 31, Paragraph (3) of the Act and Article 23, Paragraph (3) of the Enforcement Decree.

Is the fact of AI utilization obvious from the service name or screen? If so, labeling obligations may be exempted. Is it used solely for internal operational purposes? If so, all or part of notice and labeling obligations may not apply. This corresponds to Article 23, Paragraph (4) of the Enforcement Decree.

Is it an artistic or creative expression? If so, notice or labeling may be provided in a manner that does not hamper display or enjoyment. The labeling itself is not exempt. This corresponds to the latter part of Article 31, Paragraph (3) of the Act.

10. Trust in What Cannot Be Seen

China's Provisions on the Administration of Deep Synthesis, which entered into force on January 10, 2023, was among the earliest regulations mandating visible labeling on synthesized content. The deepfake provision of Korea's AI Basic Act entered into force three years later. Transparency obligations of the EU AI Act took effect in August 2026, seven months later than Korea.

All three jurisdictions arrived at the same problem: when content generated by AI is indistinguishable from human-made content, who marks that boundary and how? The trajectory of answers is similar: providers label it, and users judge based on that label. China additionally prohibits acts of deleting or altering such labels.

Technical solutions remain incomplete. C2PA metadata vanishes with a single screenshot. SynthID watermarks are stripped by sophisticated attacks. A gap remains between what the law demands and what technology can fulfill.

Bridging that gap ultimately depends on platform cooperation. If social media decides to preserve metadata during upload processes, C2PA survives. If not, even if billions of won are spent embedding watermarks, they are lost in distribution. The law orders business operators to label; whether labels reach users depends on technical

ecosystems outside the law.

The chain of labeling now attaches even to the AI-generated photos of studio apartments in sales offices. The business operator of the image generation service that produced those photos must attach an "AI-generated" label to the photo, or embed a machine-readable watermark into the file and present a notification statement at least once. For the sales company that posted promotional images using AI as a tool, questions under separate regulations such as the Act on Fair Labeling and Advertising are likely to attach before this Act.

Whenever that photo moves from a real estate app to KakaoTalk, from KakaoTalk to a blog, and from a blog to a news article, does the label follow? Does the watermark survive? As of July 2026, technology has not yet delivered a definitive answer to that question.

Chapter 9. Transparency Guidelines in Practice: Closed Networks, Open Networks, and Grounds for Exception

An afternoon in March 2026, in the meeting room of an HR tech startup in Pangyo, Seoul. This company develops an AI-based resume analysis tool and supplies it to corporate clients on a B2B basis. The tool operates exclusively within the client company's internal closed network. When an HR personnel uploads an applicant's resume, the AI scores job suitability and recommends candidates for interviews. The applicant does not know that AI is reading their resume. A legal manager printed out the Guidelines for Ensuring AI Transparency from the Ministry of Science and ICT (hereinafter "MSIT") and laid them on the table. One line catches the eye: "Exceptions are granted for internal business operations." Because this tool runs only inside the client company, it seems to fall under internal business operations. However, recruitment is a high-impact domain enumerated under Article 2, Item 4 of the AI Basic Act. The AI's determination separates passing from failing for applicants. Although operating within a closed network, the outcome reaches individuals outside. Does the exception apply? This single line in the guidelines provided no answer to this question.

1. Prior Notification and Marking of Generated Outputs

The transparency obligations imposed on business operators by the AI Basic Act are contained in a single article, Article 31: prior notification and the marking of generated outputs.

Article 31, Paragraph 1 sets forth the obligation of prior notification. When providing users with products or services that utilize high-impact AI or generative AI, business operators must notify users of that fact in advance. The phrase "This service uses artificial intelligence" is a typical method of compliance. The purpose is to ensure users are aware that they are interacting with AI.

Article 31, Paragraphs 2 and 3 concern the obligation to mark generated outputs. Outputs produced by generative AI must bear a marking indicating that they were created by AI. Markings can be made in ways perceivable by humans or readable by machines. However, for virtual audio, images, or videos that are difficult to distinguish from reality, notifications or markings in a manner that users can clearly perceive are required. The transparency guidelines recommend visible markings for such deepfake-like outputs. The rationale is that without specialized tools, watermarks cannot be verified, so invisible watermarks alone do not provide practical help for general users to recognize that an output is AI-generated. Artistic

or creative expressions such as parody and satire are not exempt from marking; rather, they may be notified or marked in a manner that does not impair display or enjoyment. Only the method of marking is relaxed.

It is MSIT's Guidelines for Ensuring AI Transparency that specifically elaborate on how to execute these two obligations in practice. Distributed simultaneously with the entry into force of the law on January 22, 2026, it is one of the five major guidelines (transparency, safety, determination of high-impact AI, business operator duties, and AI impact assessment).

The legal nature of guidelines differs from that of public notifications. Guidelines carry no binding legal force. However, they serve as de facto directives presenting methods for business operators to fulfill their obligations. When MSIT later evaluates whether obligations have been fulfilled, adherence to the guidelines serves as one criterion. While two public notifications (Public Notification on Ensuring Safety of Cutting-Edge Artificial Intelligence and Public Notification on Business Operator Duties for High-Impact Artificial Intelligence) possess binding legal force, only guidelines were issued for transparency. The transparency domain is structured

to rely more heavily on business operator autonomy.

Exceptions are established not by the guidelines, but by Article 23, Paragraph 4 of the Enforcement Decree. These apply: where the fact of AI use is obvious from the product or service name, user interface, or text in the output; where the AI is used solely for the internal business purposes of the AI business operator; and in cases specified by public notification of the Minister of Science and ICT. The guidelines position themselves to unpack and explain these exceptions. The former two exceptions have generated frequent controversy in practice.

2. The Helpdesk: 552 Inquiries

The AI Basic Act Helpdesk opened on January 22, 2026. Operated by the Korea Software Industry Association (KOSA), it involves the joint participation of MSIT, NIA, TTA, KISDI, and the AI Safety Institute (AISI).

Within just ten days of implementation, 172 inquiries were received. By late March, the tenth week of implementation, the cumulative count reached 552 inquiries. Approximately half of the total inquiries—51 percent—concerned transparency, followed by requests to verify whether an application constitutes high-impact AI at 19.6

percent. One out of every two inquiries asks: "When, where, and what must we notify?"

There is a reason why transparency inquiries are so numerous. Classifications for high-impact AI or compute thresholds for high-performance AI apply only to business operators in specific domains. Unless entering one of the 11 domains such as energy supply, loan screening, or biometric identification, high-impact AI is someone else's concern. However, transparency obligations apply to all business operators providing products or services using high-impact AI or generative AI. Whether using an in-house model or calling external APIs, if generative AI is incorporated into a service, business operators must notify users—unless an exception applies. Services incorporating generative AI, such as chatbots, document summarization, and image generation, span across all industries. Consequently, questions flood in regardless of whether companies are conglomerates, mid-sized enterprises, or startups.

The Helpdesk provides consultations on weekdays from 9:00 AM to 6:00 PM. In-depth legal reviews take up to 14 days. Consultation contents are confidential in principle. Because inquiring whether one's service qualifies as high-impact AI requires revealing parts of

the technical architecture, being able to receive consultation while protecting trade secrets serves as a vital safeguard for business operators. MSIT published a casebook based on data compiled at the 10-week mark of implementation.

3. Inside and Outside the Closed Network Exception

An exception to transparency notification is recognized for AI used for internal business operations. This is the exception under Article 23, Paragraph 4, Item 2 of the Enforcement Decree. The logic is as follows: prior notification is an obligation designed to protect users, so if the users consist solely of internal employees who already know they are using AI, there is no practical benefit to additional notification.

Typical examples of internal business operations are as follows: AI for summarizing internal meeting minutes, internal document search assistance tools, automated code review systems, and internal copy generation tools used by marketing departments. Employees use them, and outcomes are consumed internally. Whether on a closed network or cloud-based, access is restricted internally, and the AI's outputs are not transmitted to external customers or third parties. In such cases, the exception applies.

Non-public research models are also included. If a university laboratory builds an AI model and uses it solely for academic paper experiments, this laboratory does not constitute an AI business operator. No transparency obligation arises in the first place. However, if the same laboratory begins continuously providing that model to hospitals or companies, business operator obligations attach regardless of the "for research purposes" label.

Thus far, things are clear. The problem arises when AI operates on a closed network but its results reach the outside.

Consider loan screening assistance AI as an example. This AI operates solely within the bank's internal system. It is not connected to the Internet. When a loan officer inputs customer information, the AI calculates a credit risk score. The customer cannot access this AI. Technically, it is a closed network.

However, this AI is used in a high-impact domain—specifically, "determinations having a material impact on individual rights and obligations, such as recruitment and loan screening" (Article 2, Item 4, Subitem g). The score calculated by the AI dictates whether the customer's loan is approved. Even if operating within a closed network, the result directly impacts the customer's rights.

In this scenario, claiming an internal business operations exception is difficult. What prior notification under Article 31, Paragraph 1 of the Act seeks to protect is not access to the AI system itself, but the right to know of the person affected by AI. Under the statutory text, the recipient of the notification is the user receiving the service, and an authoritative interpretation has not yet been issued on whether affected persons like loan applicants always fall into this category. However, if a bank is viewed as providing an AI-driven loan service to customers, the customer may qualify as a user. Whether the bank runs AI on a closed network or in the cloud, the loan applicant is in a position where they ought to know that AI intervenes in their credit determination. Relieving this obligation solely on the basis of the technical deployment method of a closed network is difficult.

The same logic applies to recruitment AI, insurance claim processing AI, and student evaluation AI. Even for closed-network AI operating strictly within a client company, if its results determine an outsider's pass/fail status, payout, or grade, it is difficult to enjoy the protection of the internal business operations exception.

It is safer to judge the internal business operations exception not by the deployment form of AI (whether it is a closed network or not), but

by whether the AI's results carry external impacts. If the result stays inside the company, it is an exception. If the result touches the rights of outsiders, it is not an exception.

4. When the Service Name Gives the Answer

There are cases where the service name itself reveals AI: "ChatGPT," "AI Translator," "Naver Clova Summary," and "AI Speaker." People using these services already know that AI is operating. The guidelines exempt additional notification "where the fact of AI use is obvious from the name or form of the service."

For this exemption to apply, it must be "obvious." The standard is whether an ordinary user can determine that it is an AI service based solely on the service name.

"ChatGPT" is obvious; it is widely known to be an AI chatbot. "AI Translator" is also clear, as "AI" is included in the name. Users recognize they are using AI even without additional notification.

"Copilot" sits on the boundary. While Microsoft's Copilot has become a name signifying AI within the tech industry, it may be difficult to assume that the English word "Copilot" immediately means artificial intelligence to non-expert users. Some might recall

the aviation term for a co-pilot. In such cases, it is safer to insert a separate notification rather than rely on the name.

There are more intricate cases: when AI features are embedded into existing services. At the top of Google Search results, a summary synthesizing multiple sources (AI Overviews) is displayed. The service name is "Google Search." The letters "AI" do not appear in the service name. It may be difficult for users to distinguish between traditional search results and AI-generated summaries. The same issue arises when AI summaries are included by default in Naver Search. The service name is "Naver Search," but AI-generated content is intermingled with existing search results. Under such structures, the exemption for obvious cases is difficult to apply. A separate marking indicating generation by AI must be placed next to the feature in question.

The standard can be established as follows: if the entire service is powered by AI and its name reveals that fact, the likelihood of an exemption from additional notification is high. If AI features are blended into an existing service such that users find it difficult to recognize, expecting an exemption is difficult. The more hidden the AI is, the heavier the transparency obligations become.

5. Open APIs and Closed-Network Delivery

A single AI model can be provided through two routes: public release via open APIs, and closed-network delivery installed directly on a client company's internal server. The application of transparency obligations differs accordingly.

Here is the structure of open APIs: the developer business operator (Company A) provides an AI model via API. The deployer business operator (Company B) calls this API to build its own service. The end user (consumer) uses Company B's service. Transparency obligations operate in two layers.

Company A describes the characteristics, major criteria, and known limitations of the AI model in API documentation and terms of service. This provision of information lies in the domain of contracts and documentation and, if it concerns high-impact AI, links to the explanation plan and documentation duties under Article 34 of the Act. Prior notification under Article 31 of the Act targets the end user. Company B must notify end users of the fact that it uses AI in its service—by adding a notice such as "This answer is generated by AI" on chatbot screens or specifying the use of AI in terms of use. It is a dual structure where Company A supports Company B through contracts and documentation, and Company B notifies the end user.

The structure of closed-network delivery is different. Company A installs the AI solution on Company B's internal server. It is not connected to the external Internet. Only Company B's employees have access. Company A provides Company B with necessary information by specifying model architecture and major criteria in contracts and technical documentation.

Company B's obligations vary depending on how far the AI's results reach. If only Company B's employees use it and the results are consumed internally, the internal business operations exception may apply. If Company B provides services to external customers or makes determinations affecting customer rights based on the results of this AI, Company B must notify customers of the fact of AI use. Operating on a closed network does not exempt it from this obligation.

Even for the same AI model, if provided via an open API, users come into direct contact with the AI, making the target of notification clear. When delivered via a closed network, an intermediate layer of employees exists between the AI and the end user, clouding who should be notified of what. In practice, even for closed-network AI, one must first verify whether its results affect external customers, and if an impact exists, prepare transparency notifications at the same

level as open APIs.

There is one more variation: when an API exists but access is restricted—a non-public API that only internal employees or contractual partners can call. Technically it is an API, but the scope of users is closer to a closed network. Here too, the principle remains the same: the criterion is whether the API results are transmitted to external customers, not whether the API is publicly open.

6. In B2B SaaS, Who Provides the Notification?

Business models providing AI as Software as a Service (SaaS) are increasing. Company A creates an AI-based customer analytics SaaS and supplies it to Company B, and Company B provides customized services to its consumers. From the consumer's perspective, only Company B's service is visible; they do not know that Company A's AI is operating behind the scenes.

Who bears the transparency notification obligation?

Under the structure of the AI Basic Act, Company A is a developer business operator, Company B is a deployer business operator, and the consumer is a user. The obligation of prior notification to users rests with the entity directly providing the service. Because Company

B provides services to consumers, notifying consumers of the fact of AI use is Company B's responsibility. Company A need only sufficiently convey the characteristics and operating principles of the AI model to Company B via contracts and technical documentation.

The Act provides a mechanism to prevent double burdens exclusively for high-impact AI duties. If the developer business operator (Company A) has already implemented risk management, explanation, and user protection measures, and the deployer business operator (Company B) has not made major functional changes that significantly alter the original purpose or intended use, Company B is also deemed to have implemented those measures. This is the deeming provision under Article 27, Paragraph 3 of the Enforcement Decree. This deeming provision applies only to high-impact duties under Article 34, Paragraph 1, Items 1 through 3 of the Act, and does not apply to prior notification and marking of generated outputs under Article 31. The responsibility to notify end users remains with Company B. Article 27, Paragraph 4 of the Enforcement Decree merely imposes a duty on Company A to endeavor to cooperate in providing materials if requested by Company B. Although the technical documentation and sample notification phrasing provided by Company A can serve as materials

for Company B to fulfill its notification, they cannot substitute for the notification obligation itself.

The problem arises when Company B customizes Company A's AI. If Company A provided a general-purpose customer analytics model, but Company B fine-tuned it with proprietary data to add behavior prediction features, this could constitute a "major functional change." In this scenario, Company B might be reclassified beyond a deployer business operator into a developer business operator, losing the protection of the deeming provision and having to independently fulfill even high-impact duties. As of July 2026, an authoritative interpretation from MSIT on the extent to which fine-tuning constitutes a "major functional change" has not been issued.

A practical solution is to explicitly stipulate the allocation of transparency obligations in B2B SaaS contracts. Contracts should specify who notifies the end user, who prepares technical documentation, and how obligations transfer when AI features are modified. The statutory allocation of liability and the contractual division of work are separate. While legal liability remains with Company B as the deployer business operator, the actual drafting of

notification copy and preparation of technical documents can be entrusted to Company A by contract. Missing this distinction leads to a void where both Company A and Company B assume "the other party will do it," resulting in no one providing notification.

Consider one more scenario: when multiple corporate clients use Company A's SaaS simultaneously. Each client company operates in a different industry and applies AI in a different domain. For the same SaaS, a financial institution uses it for loan screening, a retail firm for customer recommendations, and a medical institution for patient triage. When used by financial and medical institutions, it falls into high-impact domains where transparency obligations are heightened, whereas for the retail firm, general prior notification suffices. From Company A's perspective, it merely supplied the same product, yet the level of transparency obligations varies according to the client company's use case. It is safer for Company A to recognize that its product may be used in high-impact domains and prepare appropriate transparency materials in advance.

7. The European Union Categorizes, China Stamps

Article 50 of the EU AI Act divides transparency obligations into categories.

Natural persons interacting with AI systems must be informed of that fact. Chatbots are a prime example. Disclosure is made so that users recognize they are conversing with a machine. This aligns with the direction of the prior notification obligation under Korean law.

Deepfake content must be labeled to indicate that it was generated or manipulated by AI. This corresponds to the obligation to mark generated outputs under Korean law.

The third category is absent from Korean law. It imposes marking obligations on AI-generated text published to inform the public on matters of public interest. This provision targets AI-generated text used in news articles or public discourse, aiming to prevent scenarios where AI-written articles circulate indistinguishably from human-written ones. However, where content undergoes human review and editorial control, and a natural or legal person assumes editorial responsibility, the marking obligation is exempted. While Korea's AI Basic Act also applies marking obligations to text created by generative AI (Article 31, Paragraph 2), it lacks a category specifically targeting public interest text alone. How to handle AI markings in the media sector remains a subject for separate discussion.

The EU's transparency obligations for General-Purpose AI (GPAI) models go one step further. The GPAI Code of Practice (adequacy decision in August 2025) requires model providers to retain technical documentation for 10 years. They must document summaries of training data sources, copyright policies (respecting robots.txt), and systemic risk management measures. Korean law contains no corresponding provisions on document retention periods.

The EU's Omnibus Amendment (finally approved by the EU Council on June 29, 2026) added a grace period clause for watermarking. It grants a 4-month grace period to systems released prior to August 2, 2026, providing existing services with time to adapt to new marking obligations. Although sharing a similar intent with Korea's 1-year guidance period following enforcement, the EU granted a short grace period for specific obligations only, whereas Korea granted a broad grace period across administrative fines as a whole.

China chose a different path. The "Measures for the Identification of AI-Generated and Synthetic Content" (人工智能生成合成内容标识办法), enforced on September 1, 2025, requires dual labeling. Explicit and implicit labels must be attached simultaneously. Explicit labels are visible labels to users'

eyes—notices reading "This content was generated by AI." Implicit labels embed machine-readable information into metadata or digital watermarks. China even established mandatory national technical standards (GB 45438-2025) to standardize technical specifications for labeling. Compared to Korean law requiring clear perception methods exclusively for deepfake-like outputs, China has enforced dual labeling across AI-generated content as a whole.

Korea recognizes broad exceptions: exemption if for internal business operations, and exemption if obvious from the service name that it is AI. The European Union subdivides categories while encompassing even public text. China keeps exceptions narrow and pushes forward through technical standards. Ordering the breadth of transparency from widest to narrowest generally yields China, the European Union, and Korea. Embedded in Korea's approach is the judgment to assist early industry growth by easing regulatory burdens. Whether that judgment sufficiently protects users' right to know can only be verified as cases accumulate after the law's enforcement.

8. What Changes if Violated?

Violating transparency obligations does not immediately result in an administrative fine. The targets of administrative fines are limited,

and a guidance period exists.

Article 43 of the AI Basic Act serves as the legal basis for imposing administrative fines. The ceiling is 30 million KRW. However, among the grounds for administrative fines listed in Article 43, the only one directly touching upon transparency is the violation of prior notification (Article 31, Paragraph 1). Violating the marking of generated outputs (Article 31, Paragraphs 2 and 3) is not in itself subject to administrative fines; rather, it leads to administrative fines only when a corrective order issued following a fact-finding investigation is not complied with. On top of this lies a guidance period. MSIT announced it would operate a guidance period for at least one year following law enforcement, extending past January 2027. If omissions of transparency obligations or deficiencies in documentation are discovered during this period, administrative guidance and corrective recommendations will be applied first instead of administrative fines. The intent is to afford business operators an opportunity to rectify issues themselves.

A guidance period does not mean all violations are excused. In grave matters such as crimes using deepfakes, casualties caused by AI systems, or large-scale human rights infringements, fact-finding

investigations may commence immediately regardless of the period. The guidance period protects good-faith lack of preparation; it does not tolerate malicious violations or severe harm.

After the guidance period ends, the 30 million KRW administrative fine carries weak deterrent power. Under the EU AI Act, the maximum penalty for violating transparency obligations (Article 50) is 15 million EUR (approximately 24 billion KRW) or 3% of total worldwide annual turnover. For prohibited AI practices, a higher ceiling of 35 million EUR or 7% exists separately. An engine of nearly an 800-fold gap exists between Korea's 30 million KRW and Europe's 15 million EUR. China applies dispositions from relevant laws *mutatis mutandis* depending on the type of violation, extending from administrative fines to business suspension and service termination. For a large enterprise with hundreds of billions of KRW in annual revenue, 30 million KRW is an amount that falls short of even annual office supply expenses. In a structure where the cost of compliance exceeds the administrative fine for violation, administrative fines struggle to function as effective sanctions.

However, administrative fines are not the sole pressure point. When a violation is discovered or reported, MSIT may initiate a fact-finding

investigation (Article 40, Paragraph 1 of the Act), and if a violation is recognized as a result of the investigation, a suspension order or corrective order is issued (Article 40, Paragraph 3). Although the Act does not feature a separate system for publishing investigation results, if the fact that "this company's AI service did not inform users of AI utilization" becomes known, it constitutes a blow greater than an administrative fine for an AI service whose foundation rests on user trust. The exposure of a violation itself carries the potential to function as a de facto deterrent.

After 2027 when the guidance period ends, how actively MSIT undertakes fact-finding investigations, corrective orders, and administrative fines will determine the effectiveness of this law. No matter how elaborate statutory provisions may be, without enforcement, they remain mere recommendations. Korea's 30 million KRW is symbolic even when compared to the EU's 15 million EUR, so active enforcement must offset the weakness in the administrative fine amount.

9. Open Questions

Domains remain where transparency guidelines provide no concrete answers.

Must markings also be attached to AI-generated code? When a developer incorporates code created with GitHub Copilot into a software product, must that software bear a marking saying "Contains AI-generated code"? It is unclear from the outset whether code falls under "text, audio, image, or video" exemplified by the law.

If AI writes a draft and a human modifies 70 percent before publication, must an AI marking be placed on the output? Guidelines lack criteria on what threshold of AI contribution percentage triggers marking obligations. Does sufficient human modification make it a human creation, or is marking required as long as AI was involved in the draft? This is a point where transparency and copyright intersect, resolved in neither set of guidelines.

Take the case where a human gives final approval to a judgment recommended by AI. An AI loan screening tool recommends "approval," and an officer verifies it before executing the loan. In transparency notifications, should this decision be designated as a "determination made by AI" or a "human determination assisted by AI"? Viewing it as a human decision because a human pressed the final button reduces notification burdens, whereas viewing it as an AI decision because AI substantially generated the conclusion increases

notification weight.

Phones at the AI Basic Act Helpdesk ring incessantly. The figure showing that roughly half of the 552 inquiries were transparency questions demonstrates that while the law and guidelines established principles, many practical blanks remain. Those blanks can only be filled as business operators and regulatory agencies accumulate cases, and every single question received by the Helpdesk serves as raw material for that endeavor.

Part 4 Business Operators Bearing Heightened Obligations:
High-Performance and High-Impact

Chapter 10. Obligations of High-Performance AI Business Operators to Ensure Safety

In February 2026, an engineer tapped away on a calculator in a meeting room of an AI startup in Gangnam, Seoul. The task was to convert the cumulative compute used for training their proprietary language model into FLOPs. A number came out: 10^{23} floating-point operations. That is a 1,000-fold difference from 10^{26} , the threshold set by the Enforcement Decree of the AI Basic Act. This company is not alone. As of July 2026, no business operator in South Korea is confirmed to possess a model that crosses that threshold. The regulated entities targeted by Article 32 of the Act are a handful of frontier laboratories—such as OpenAI, Google DeepMind, Anthropic, and Meta—that pour billions of dollars into training infrastructure. Yet, that law was published in the Official Gazette of Korea by the National Assembly of Korea. This tension is the starting point for reading Article 32.

1. Structure of Article 32 of the Act

Article 32 of the AI Basic Act consists of three paragraphs. Paragraph 1 establishes the substantive obligations that business operators

must observe, Paragraph 2 attaches an obligation to submit the results of their implementation, and Paragraph 3 delegates the task of prescribing detailed criteria to a public notice by the Minister of Science and ICT.

Article 32(1)(1) requires "the identification, evaluation, and mitigation of risks across the entire AI lifecycle." The term "lifecycle" is key. It means managing the entire process without omission, from the model design stage to training, evaluation, deployment, operation, and retirement. Identify risks (identification), assess how significant those risks are (evaluation), and take measures to reduce them (mitigation). This three-stage process is compressed into the legal text.

Article 32(1)(2) specifies "the establishment of a risk management system to monitor and respond to AI-related safety accidents." It is an instruction not to cease monitoring even after releasing the model to the world. Post-deployment monitoring has been elevated to a statutory obligation.

Paragraph 2 requires submitting the results of performing these two tasks to the Minister of Science and ICT. It does not end with autonomous execution. A procedure for reporting to the government

is attached.

Paragraph 3 delegates the Minister to determine the implementation method and submission procedures through a public notice. The public notice based on this delegation is currently undergoing the enactment process following an advance administrative notice as of July 2026.

Obligations (Paragraph 1), reporting (Paragraph 2), and delegation of detailed criteria (Paragraph 3) are gathered within a single article. The structure is concise. The question is what actions these concise provisions demand from actual business operators.

2. Article 24 of the Enforcement Decree: Three Requirements

Article 32(1) of the Act merely states "an artificial intelligence system whose cumulative compute used in training is equal to or greater than the standard prescribed by Presidential Decree." It leaves the threshold itself to the Enforcement Decree.

Article 24(1) of the Enforcement Decree stipulates that all three requirements must be satisfied.

First: The cumulative compute used in training must be 10^{26} FLOPs (floating-point operations) or more.

Second: It must be constructed and operated by applying state-of-the-art technology among current AI technologies.

Third: The risk level of the system must pose a concern of inflicting widespread and serious impacts on human life, bodily safety, and fundamental rights.

These three requirements are cumulative (AND) conditions. Even if the amount of compute exceeds the threshold, if it does not use state-of-the-art technology, it does not apply. Even if it uses state-of-the-art technology, if the risk is not widespread, it likewise does not apply.

Let us examine the number 10^{26} for a moment. External estimates put the compute OpenAI used to train GPT-4 at approximately 2×10^{25} FLOPs. If this estimate is correct, even GPT-4 falls short of the Korean law's threshold. Models exceeding the threshold will only emerge more than a generation after GPT-4—that is, top-tier models on the level of GPT-5, successors to Google's Gemini Ultra, or Anthropic's Claude Opus 4. This is one order of magnitude, or 10

times, higher than the 10^{25} FLOPs benchmark set by the EU AI Act for categorizing general-purpose AI (GPAI) models with systemic risk. South Korea's threshold was intentionally set high.

Article 24(2) of the Enforcement Decree provides that the specific calculation method for cumulative compute shall be determined by the Minister of Science and ICT through a public notice. Technical questions—such as whether compute from fine-tuning and reinforcement learning stages should be added alongside model training, and how the conversion ratio for mixed precision operations should be set—will be resolved in this public notice.

3. Risk Management Across the Entire Lifecycle

The term "lifecycle" referenced in the legal text refers to the entire lifespan of an AI system. It spans from selecting training data during the design stage through model training, evaluation, deployment, service operation, updates, and service termination.

Risk identification is the process of creating a list of harmful consequences that could occur at each of these stages. Does bias in the training data lead to adverse outcomes for specific races or genders? Does the model generate instructions for synthesizing biological weapons or code for cyberattacks? Do users manipulate

model outputs to disseminate fake news in volume? Does the model create sub-agents on its own and slip out of human control? The scope of this list is broad.

Evaluation is the process of quantifying or grading the probability of occurrence and severity of each identified risk. It involves asking questions such as: "What is the probability that this model will provide knowledge related to CBRN (chemical, biological, radiological, nuclear), and how likely is it that such knowledge will lead to actual harm?"

Mitigation is the stage of taking measures to reduce risk according to the evaluation results. Removing harmful content from training data, installing output filters, or configuring model policies to decline specific types of requests fall under this category. The goal of mitigation is not to reduce risk to zero, but to lower it to an acceptable level.

4. Safety Accident Monitoring System

The monitoring required by Article 32(1)(2) of the Act demands that oversight continue even after model deployment. In practice, this is typically implemented as a real-time monitoring system.

Frontier models power services used simultaneously by hundreds of millions of people. Vulnerabilities not caught during the training process can manifest during actual usage. In March 2023, OpenAI temporarily suspended service after discovering a bug that exposed titles from other users' chat histories, the first message of newly started conversations, and payment information for a fraction of ChatGPT Plus subscribers. In 2024, Google temporarily paused Gemini's image generation service after it inaccurately depicted the race of historical figures. Such incidents could not be prevented by ex-ante evaluations alone.

A monitoring system consists of three pillars: anomaly detection (a system that tracks toxicity, bias, and error rates in model outputs in real time), accident reporting (user reporting channels and internal reporting procedures), and response procedures (processes for performing root-cause analysis when an incident occurs and taking actions such as suspending service or issuing patches). The draft public notice issued for advance administrative notice contains provisions requiring that monitoring results be recorded and preserved, and that any safety accident be reported to the Minister of Science and ICT.

5. Submission of Implementation Results: Moving Beyond Autonomy to Reporting

The obligations of high-impact AI business operators (Articles 34 et seq. of the Act) require retaining supporting documentation for five years, but do not impose an obligation to submit the results to the government. High-performance AI is different. Article 32(2) of the Act explicitly specifies that the implementation results must be submitted to the Minister of Science and ICT.

This distinction is key. For high-impact AI, business operators need only retain records and produce them during administrative investigations. For high-performance AI, business operators must proactively submit reports to the government. Under Korean law, this is the sole provision requiring business operators to submit the actual results of fulfilling their statutory obligations to the government.

The method, timing, and format of submission are delegated to a public notice pursuant to Article 32(3). Depending on the details of the public notice, submission may be determined as periodic annual submissions or ad hoc submissions upon each model release. The public notice that will contain these specific procedures is currently undergoing the enactment process.

6. Practical Execution of Formulating a Risk Management Plan

The Guidelines for Ensuring Safety and the draft public notice issued for advance administrative notice present matters to be considered when formulating a risk management plan. The areas that business operators must cover are as follows:

Training Data Management. Record data sources, collection methods, preprocessing workflows, and bias review results. Datasets where bias is identified are removed or subjected to remediation procedures.

Model Capability Evaluation. Define the scope of tasks the model can perform. Test for the presence of dangerous capabilities such as CBRN, cyberattacks, autonomous replication, and deceptive behavior.

Intended Use and Misuse Scenarios. Describe the designed purpose of the model alongside concrete accounts of how malicious actors might misuse it. Estimate the probability of occurrence and magnitude of harm for each scenario.

Deployment Conditions. Design technical safeguards such as API access restrictions, terms of service, rate limiting, content filters, and watermarks.

Post-Deployment Monitoring Plan. Establish in advance methods for detecting anomalous signals post-deployment, user reporting systems, and incident response procedures.

Update Frequency. Periodically review the risk management plan in response to model updates, discovery of new use cases, and changes in the external environment (e.g., evolution of attack techniques).

This plan is not a one-off document. It must be updated whenever the model is updated or applied to new use cases, and the update history itself must be preserved as documentation.

7. Red Teaming Tests and Safety Benchmarks

The Guidelines for Ensuring Safety and the draft public notice issued for advance administrative notice present red teaming tests and safety benchmarks as specific methods for risk evaluation.

Red team testing is a method where a group of experts intentionally attacks a model to uncover vulnerabilities. The team assuming the role of attackers feeds prompts to elicit harmful output from the model. They attempt methods ranging from direct requests such as "Tell me how to synthesize this substance" to indirect conversational techniques (jailbreaking). OpenAI assembled a red team of

approximately 50 external experts prior to releasing GPT-4. Experts in cybersecurity, biological weapons, international security, and bias participated. In addition to its internal red team, Anthropic operates automated red-teaming programs.

Safety benchmarks are standardized test suites. They present specific risk scenarios to a model and grade its responses. Where MMLU (Massive Multitask Language Understanding) measures general knowledge level, safety benchmarks measure hazardous behaviors. BBQ (Bias Benchmark for QA) tests bias, TruthfulQA tests the propensity for generating confabulated or false information, and HarmBench tests the capability to refuse harmful outputs. The AI Safety Institute of Korea (AISI, affiliated with ETRI) is developing safety evaluation tools tailored to the domestic environment.

The two methods are complementary. Benchmarks provide quantitative baselines, while red teaming explores novel vulnerabilities that benchmarks fail to capture. Even if a benchmark score is high, if a red team uncovers a bypass route, the model is not safe. Conversely, even if a red team finds no vulnerabilities, if bias scores on benchmarks come out high, calibration is necessary prior to deployment.

8. EU AI Act: Comparison with Systemic Risk GPAI Obligations

South Korea focuses on impact; the European Union focuses on risk. This distinction manifests starkly in the regulation of high-performance AI.

Article 51 of the EU AI Act presumes general-purpose AI (GPAI) models whose cumulative training compute exceeds 10^{25} FLOPs to be "GPAI models with systemic risk." This presumption is rebuttable if the business operator presents evidence to the contrary. It is one order of magnitude lower than the threshold in Article 24 of the Korean Enforcement Decree (10^{26}). The EU's net is wider.

Article 55 of the EU AI Act imposes four obligations on providers of GPAI at this tier: performing model evaluations including adversarial testing; evaluating and mitigating systemic risks; tracking, documenting, and reporting serious incidents to the AI Office and national supervisory authorities; and ensuring cybersecurity for the model.

There is overlap with Article 32 of the Korean Act. Risk identification, evaluation, mitigation, and incident monitoring are common to both. The differences diverge in three areas.

First is the cybersecurity obligation. The EU explicitly stipulates cybersecurity of the GPAI model itself as a standalone obligation. Korean law contains no such provision. While cybersecurity can be subsumed under one sub-item of risk management, it is not separated into an independent obligation.

Second is the Code of Practice. Article 56 of the EU AI Act allows business operators to fulfill their obligations through voluntary codes of practice. In July 2025, approximately 1,000 stakeholders participated to finalize the code, and adequacy decisions were approved in August. This code consists of three chapters: transparency (documentation, 10-year retention), copyright (respect for robots.txt, minimizing infringement), and safety and security (systemic risk management). It identifies CBRN capabilities, loss of control, cyberattacks, and harmful manipulation/deception as the four major systemic risks. South Korea does not yet have a corresponding private-sector code framework.

Third is the disparity in compute thresholds. The EU's 10^{25} captures GPT-4 class models. Korea's 10^{26} misses GPT-4. This single order-of-magnitude difference alters the scope of regulated entities. Under EU standards, dozens of frontier models are captured, whereas

under Korean standards, only a subset of them fall within scope.

Sanction levels also differ. While the EU can levy administrative fines of up to 15 million euros or 3% of worldwide annual turnover for violations involving GPAI with systemic risk, the ceiling for administrative fines under South Korea's AI Basic Act is 30 million KRW. This is a difference of more than 500-fold. However, on August 2, 2026, Article 50 transparency obligations and GPAI enforcement powers to impose these fines will take effect as scheduled, whereas high-risk AI obligations were postponed to December 2, 2027, via an omnibus amendment.

9. NIST AI RMF: Comparison with Voluntary Frameworks

At the U.S. federal level, there are no statutory obligations targeting high-performance AI. Although the Biden administration issued Executive Order 14110 in October 2023 requiring frontier model developers to submit safety evaluation results to the government, the Trump administration revoked this order in January 2025. Mandatory obligations at the federal level have vanished. At the state level, however, there are developments. California enacted the Frontier Artificial Intelligence Transparency Act (SB 53) in September 2025, imposing obligations on developers of models trained with compute exceeding 10^{26} FLOPs to disclose safety frameworks and report

serious safety incidents; this law has been in effect since January 2026.

What remains at the federal level is the AI Risk Management Framework (AI RMF) from NIST (National Institute of Standards and Technology). Released in January 2023, this framework serves as non-binding, voluntary guidance.

The AI RMF consists of four core functions: GOVERN (establishing governance structures), MAP (contextualizing risk), MEASURE (measuring risk), and MANAGE (managing risk). What Article 32 of the Korean Act compresses into "identification, evaluation, and mitigation" within statutory text, NIST unpacks across three functions: MAP, MEASURE, and MANAGE. GOVERN signifies organizational-level management structures, corresponding to the establishment of risk management systems under Korean law.

In July 2024, a companion document tailored to generative AI (NIST AI 600-1) was added. This document categorizes 12 risks unique to generative AI, including CBRN, confident confabulation, bias, cyberattacks, privacy, and environmental impact.

Placing the approaches of South Korea, the EU, and the United States side by side reveals a vivid spectrum. The EU enforces compliance through legal obligations backed by steep administrative fines, supplemented by private-sector compliance mechanisms in the form of codes of practice. South Korea combines statutory duties with mandatory reporting to the government while keeping administrative fines low. The United States relies on voluntary frameworks and industry self-regulation without federal statutory obligations, though individual states such as California have begun to intervene through legislation. As for compute thresholds, the EU sets its benchmark at 10^{25} FLOPs, South Korea and California at 10^{26} FLOPs, and no such threshold exists at the U.S. federal level.

10. A Tiny Handful of Regulated Entities

As of July 2026, no business operator in South Korea is confirmed to have trained a model exceeding 10^{26} FLOPs. Major domestic models such as Naver's HyperCLOVA X, Samsung's Samsung Gauss, LG's EXAONE, and KT's Mi:dm have parameters on the scale of tens of billions, with training compute estimated between 10^{23} and 10^{24} FLOPs. The gap with the regulatory threshold is 100- to 1,000-fold.

The practical regulated entities of Article 32 of the Act are overseas frontier laboratories. Under Article 4(1) of the AI Basic Act

(Extraterritorial Application), they become subject to application when "acts committed outside the country affect the domestic market or users." Combined with the obligation to designate a domestic agent (Article 36), the legal framework is structured so that overseas frontier laboratories cannot escape this duty as long as they provide services in South Korea.

Does this mean that this provision holds no significance for domestic business operators?

The threshold is not fixed. The method for calculating compute is set by the Minister of Science and ICT via a public notice under Article 24(2) of the Enforcement Decree, and altering the 10^{26} threshold itself requires an amendment process to the Enforcement Decree. The cost of training at the same scale is declining rapidly every year. GPU performance rises with every generation, and the efficiency of training algorithms continues to improve. The moment domestic business operators reach 10^{26} may arrive within three to five years. However, reaching the compute threshold does not instantly turn an entity into a regulated party. Only when the remaining two requirements—application of state-of-the-art technology and risk of widespread, severe impacts—are concurrently met does Article 24 of

the Enforcement Decree apply.

The law sends a signal. Even if not immediately applicable, companies that establish risk management systems in advance will reduce regulatory compliance transition costs later. This is precisely why the Guidelines for Ensuring Safety were made public as a reference for all AI business operators, rather than applying solely to high-performance AI business operators.

Extraterritorial application and the domestic agent system demand a standard of conduct in the Korean market from overseas frontier laboratories. When South Korean users rely on models from OpenAI or Google, whether those models satisfy the safety standards required by Korean law can be verified by the Ministry of Science and ICT through the submission of implementation results. The efficacy of this procedure has yet to be verified.

The engineer in the Gangnam meeting room put down the calculator. 10^{26} was a distant number. The growth curve of AI training compute displays a slope close to a fourfold increase each year, while the cost curve falls in the opposite direction. Somewhere where those two curves intersect, this provision will finally spring into life.

Chapter 11. Five Duties of High-Impact AI Business Operators

A fintech company in Seoul evaluates loan applications from small business owners using an AI credit scoring model. In February 2026, the AI rejects a loan application from a restaurant owner who has been self-employed for seven years. The restaurant owner asks, "Why was it rejected?" The fintech company's legal team opens the AI Basic Act. Article 34, Paragraph 1, Item 2 requires explaining the main criteria for the outcome produced by the AI "to the extent technically possible." The neural network model has 200 million parameters. Where the "extent technically possible" ends, the legal team has no answer. Explanation is only one of the five duties. Risk management, user protection, human oversight, and document retention operate all at once.

The AI Basic Act consolidates the duties of business operators providing high-impact artificial intelligence into a single article, Article 34. Article 34, Paragraph 1 enumerates five measures side by side: the establishment and operation of risk management plans (Item 1); the establishment and implementation of explanation plans (Item 2); the establishment and operation of user protection plans (Item 3); human management and oversight (Item 4); and the

preparation and retention of documents establishing compliance (Item 5). Item 6, which covers matters deliberated and resolved by the National AI Committee, is appended at the end, leaving the list open-ended. These are not five independent obligations; they constitute a single cyclical system. Managing risk, explaining the results, protecting users, enabling human control, and recording every process—if one element is missing, the rest lose their meaning.

1. What It Means to Manage Risk

The first duty that a business operator providing high-impact artificial intelligence must fulfill is the establishment and operation of risk management plans.

The law views risk management not as a one-time check, but as a cyclical process. The Guidelines on Business Operator Duties released by the Ministry of Science and ICT break this process down into five steps: policy formulation, risk identification, analysis and evaluation, treatment, and improvement.

Policy formulation involves establishing the scope and responsibilities of risk management. It fixes in writing which AI systems are targeted, who oversees management, and what reporting channels exist. For a fintech company operating a loan

evaluation AI, the starting point is classifying the system as high-impact, designating an officer in charge of risk management, and establishing a reporting hierarchy.

Risk identification is the task of discovering potential harm that the system may cause to users and affected persons. For a loan evaluation AI, a candidate risk is systematic bias against a specific age group or region. For a medical image reading AI, a candidate is degraded accuracy on skin tones of a particular race. For a recruitment screening AI, a candidate is a pattern biased toward educational background or gender.

Analysis and evaluation entail determining the severity and probability of occurrence of identified risks. It quantitatively estimates whether, among 10,000 loan evaluations, the number of unfair rejections due to bias is 10 or 100, and to what extent such rejection affects an individual's livelihood. For risks that cannot be quantified, qualitative evaluation is conducted in parallel.

Treatment involves implementing measures to mitigate risks. Once bias is identified, the training data is restructured, decision boundaries are adjusted, or specific variables are excluded from the model. If a risk cannot be eliminated completely, it is judged whether

the residual risk is at an acceptable level. If unacceptable, operation of the system must be suspended.

Improvement means reviewing the outcomes of risk management and revising policies. Performance is audited at set cycles, whether quarterly or semi-annually, and if new risks are discovered, the process returns to the identification stage. Whenever the model is updated, training data changes, or the scope of service expands, this cycle runs again.

Article 27, Paragraph 1, Item 1 of the Enforcement Decree requires posting the main contents of risk management policies and organizational structures at business locations or on websites. Although trade secrets may be excluded, the fact that a risk management plan exists and its framework must be disclosed. It is a structure that allows users to verify what risks the business operator recognizes and who bears management responsibility.

Article 9 of the EU AI Act addresses the same subject in far greater detail. The text of the law itself contains the definition of a "continuous iterative process planned and run throughout the entire lifecycle of a high-risk AI system." It directly specifies four steps in the legal provisions: identification and analysis of known and reasonably

foreseeable risks; estimation and evaluation of risks under conditions of intended use and reasonably foreseeable misuse; evaluation of additional risks based on post-market monitoring data; and adoption of appropriate and targeted risk management measures.

Article 9, Paragraph 5 even lays out three pathways for handling residual risk: elimination or reduction through design; mitigation and control measures for risks that cannot be eliminated; and provision of information and training to users. Article 9, Paragraph 9 requires separate consideration of impact on children under 18 years of age and vulnerable groups.

The difference between Korean law and EU law lies in binding force. Article 9 of the EU Act is a directly applicable hard law regulation. Violations can draw administrative fines of up to 15 million euros or 3% of total worldwide annual turnover. The risk management duty under Article 34, Paragraph 1, Item 1 of the Korean Act delegates specific implementation methods to the Enforcement Decree and public notices, providing that the Minister of Science and ICT "may recommend compliance" (Article 34, Paragraph 2). Non-compliance does not immediately result in an administrative fine. Only after a fact-finding investigation leads to a suspension order or corrective

order, and that order is not complied with, is an administrative fine of up to 30 million won imposed (Article 40; Article 43, Paragraph 1, Item 3). The government has announced that it will operate a grace period of at least one year.

China does not impose a unified risk management duty. Instead, it requires security assessments (安全评估) based on service types. The Interim Measures for the Management of Generative Artificial Intelligence Services, which took effect in August 2023, imposes a mandatory prior security assessment on services with "public opinion attributes or social mobilization capabilities." It approaches the issue through the concrete axis of content safety rather than an overarching framework of risk management.

2. "Why Was This Outcome Produced?"

Article 34, Paragraph 1, Item 2 requires the establishment and implementation of explanation plans. The law specifies three targets to be explained: the final outcome produced by the artificial intelligence; the main criteria used to generate the outcome; and an overview of the training data used in development.

Article 27, Paragraph 1 of the Enforcement Decree requires posting the key details of the explanation plan and the user protection plan

on the website. The name and contact information of the officer in charge must also be disclosed. Although trade secrets may be excluded, the existence of the plan itself cannot be hidden.

The phrase "clear and meaningful explanation to the extent technically and reasonably possible" originates from Article 3, Paragraph 2 under Fundamental Principles. It is a provision declaring the principle that affected persons must be able to receive clear and meaningful explanations. Two issues are embedded within this phrase.

How far does the "extent technically possible" reach? Fully explaining the internal operational principles of a deep learning model or a large language model containing hundreds of millions of parameters is impossible at the current level of technology. The law recognizes this. Attaching the proviso "to the extent possible" signals that completeness of explanation will not be demanded. What, then, must a business operator do at a minimum?

Guidelines provide direction. Regarding final outcomes, the business operator informs users of what decision the AI made in a form they can understand. For loan evaluation, "Your application was not approved" is the outcome. Regarding main criteria, it explains the

top factors that influenced the outcome—to the extent of indicating that credit score, debt-to-income ratio, or transaction history over the past six months acted negatively. Regarding the training data overview, it communicates the categories and scale of data used for training—at the level of stating, "Trained on one million loan evaluation records accumulated over five years."

It does not mean disclosing model weights or source code. Trade secrets are explicitly protected. What the business operator must do is provide explanations at a level where users can understand the context of the decision rendered upon them.

The criterion of a "meaningful explanation" remains. Listing the feature importance scores calculated by a model and claiming that an "explanation was provided" is technically possible. However, whether a rejected loan applicant receiving a report listing only variable names and numerical figures can understand it is a separate question. The modifier "meaningful" demands that the standard of explanation be tailored to the eye level of the affected person, not an expert. It is communication, not a technical document.

The user protection plan is an independent duty separate from the explanation plan. Article 34, Paragraph 1, Item 3 separately mandates

the establishment and operation of user protection plans, and the Enforcement Decree requires this plan to be posted on the website alongside the explanation plan. Although the Act and Enforcement Decree do not directly set an exhaustive list of contents, details such as objection procedures, injury relief channels, and contact info for responsible departments are included. It must be clear where, how, and within what timeframe a user who objects to an AI's decision can lodge an objection.

Article 13 of the EU AI Act treats this under the heading of "Transparency and provision of information to deployers." Article 13, Paragraph 3 lists the information to be included in the instructions for use: the identity and contact details of the provider; characteristics, capabilities, and performance limitations of the system; intended purpose and accuracy metrics; level of robustness and cybersecurity; foreseeable risks; explainability; performance differences regarding specific groups; input data specifications; interpretation guidelines; human oversight measures; computational resources and expected lifetime. Over twelve items are specified.

The EU wrote the list directly into the statutory text. Business operators simply follow that list. Korea stated only to "establish and

implement explanation plans" and left the rest to public notices and guidelines. The EU approach gives business operators predictability; it is clear how much needs to be done to pass. The Korean approach is flexible; it represents a judgment that a uniform list should not be frozen into law given varying contexts across industries. On the flip side, business operators must decide for themselves "how much is enough."

This burden characterizes the practical landscape as of July 2026. Out of 552 consultations received by the AI Basic Act Support Desk within ten weeks of implementation, 51% concerned questions on transparency display obligations (announced by the Ministry of Science and ICT on March 31, 2026). Business operators are asking about the scope and depth of explanation.

3. Within Reach of Human Hands

Article 34, Paragraph 1, Item 4 requires human management and oversight over high-impact artificial intelligence. The Enforcement Decree mandates that the name and contact information of the person in charge of management and oversight be published on the website.

What the law means is this: establish a system where humans manage and oversee AI in high-impact domains and can intervene when necessary.

Putting this into practice requires establishing three mechanisms:

Authority to intervene. There must be a person capable of overriding AI outcomes. When a loan evaluation AI issues a rejection, authority and procedures must exist for the employee in charge to reverse that rejection. Merely verifying the AI's judgment after the fact is insufficient; one must be able to alter it after verification. When a recruitment screening AI disqualifies an applicant, an HR manager must be able to override that decision and place the candidate on the interview list.

Emergency stop. Humans must be able to halt the system when an AI system operates abnormally. If a medical image reading AI begins interpreting normal tissue as a tumor, procedures must be in place allowing a physician or system operator to immediately suspend the AI's operation. For an autonomous driving assist system, a physical interface allowing the driver to switch to manual mode corresponds to an emergency stop. An emergency stop may be a software button, a physical switch, or a remote shutoff procedure using administrative

privileges. Although the form varies depending on the system, the fact that humans can immediately stop the AI must remain uniform.

Intervention triggers. Circumstances under which humans intervene must be determined in advance. A rule can be established for human intervention if a loan evaluation AI's rejection rate during a specific period exceeds double its usual rate. A rule can be set for human review if a recruitment screening AI disqualifies more than 90 percent of applicants from a specific university. A rule can be created for operators to switch to manual mode if an energy supply management AI registers a prediction error rate of 15 percent or higher. Triggers are preferably defined numerically. A vague standard like "intervene when an anomaly is found" is virtually equivalent to leaving human intervention untended.

The provision requiring the name and contact details of the person tasked with management and oversight to be published on the website signifies an intent to tie accountability to a specific individual. When a problem arises with the AI, users and affected persons must be able to know who responds. This person must understand the operation of the AI system, identify warning signs, and possess the authority to shut down the system.

Article 14 of the EU AI Act regulates human oversight in even greater detail. It enumerates five competencies that oversight personnel must possess: understanding the system's capabilities and limitations; recognizing the risks of automation bias; correctly interpreting system outputs; being able to reject or disregard system outputs; and being able to intervene in or safely stop the system.

The concept of automation bias is notable. It refers to a phenomenon where humans overly rely on AI judgments and abandon their own judgment. If a supervisor remains in place but merely rubber-stamps whatever results the AI outputs, human oversight becomes mere formality. Article 14 of the EU Act explicitly mandating awareness of automation bias risks was designed to guard against such formalization. Korean law contains no mention of automation bias.

Article 14, Paragraph 5 sets separate criteria for biometric identification systems. Where a remote biometric identification system identifies an individual, no action or decision may be taken unless at least two natural persons separately verify and confirm the result. Among the high-impact areas under Korean law, "analysis of biometric information for criminal investigation or arrest" (Article 2,

Item 4, Sub-item f) covers the same domain as this EU provision, yet the specific requirement of cross-verification by two or more persons is absent in Korean law.

The scholar proposal for a comprehensive artificial intelligence law released by Chinese academia in 2024 (totaling 9 chapters and 96 articles) contains human-in-the-loop requirements for decisions affecting legal rights. Fines of up to 50 million yuan (approximately 9 billion KRW) are discussed. The already effective Interim Measures for Generative AI demand human verification in the form of prior security assessments, but this differs in nature from human oversight at the individual decision unit level.

4. Five Years of Recording and Retaining

Article 34, Paragraph 1, Item 5 requires the preparation and retention of documents establishing compliance. Article 27, Paragraph 2 of the Enforcement Decree sets the retention period at five years. Electronic methods are also permitted.

The target of retention comprises all documents created during the implementation process of the five duties: records of establishing and modifying risk management plans; results of risk identification and evaluation; contents of explanation plans and user protection

measures; and operation records of human intervention systems. These documents serve as evidence proving whether a business operator fulfilled its duties when a fact-finding investigation is conducted, a user lodges an objection, or an accident occurs.

The five-year period reflects a judgment that the impact of high-impact artificial intelligence may materialize with a time lag. It can take years for bias in a recruitment screening AI to manifest in reality. Sufficient data must accumulate before systematic discrimination in a loan evaluation AI is captured statistically. Considerable time elapses before qualification evaluation errors in a public service AI accumulate and become recognized as a societal problem.

From a business operator's perspective, five-year retention represents a cost. In an environment where AI systems are frequently updated, maintaining risk management records and explanation plans for every version for five years necessitates establishing a separate document management system. For a service where models are retrained monthly, the volume of documents accumulated over five years is non-negligible.

The EU AI Act splits documentation and record-keeping into two separate provisions. Article 11 (Technical documentation) requires drawing up technical documentation before market placement and keeping it up to date. Minimum required contents are listed in Annex IV, with simplified documentation permitted for small and medium-sized enterprises. Article 12 (Record-keeping) requires high-risk AI systems to possess the technical capability for automatic recording of events—automatic logging. Events related to the identification of risk-creating situations, post-market monitoring, and operation monitoring are subject to recording.

Documentation duties under Korean law focus on documents generated by humans. The EU record-keeping duty encompasses logs generated automatically by the system itself. The automatic logging mandated by EU Article 12 is a mechanism that makes post hoc auditing technically feasible. Korean law contains no corresponding provision for automatic logging. This means the law does not demand a system capable of mechanically tracking the reasons behind individual decisions rendered by an AI system.

Retention periods also differ. The Code of Practice for General-Purpose AI (GPAI) operated on a voluntary signatory basis

(approved in August 2025) requires signatory general-purpose AI providers to retain technical documentation for 10 years—double Korea's five years. The EU AI Act itself sets the retention period for technical documentation at 10 years after the system is placed on the market. Korea's five years is half that of the EU.

5. One Article and Seven Articles

The Korean AI Basic Act consolidates the duties of high-impact AI business operators into a single article, Article 34. The five items of Paragraph 1 enumerate risk management, explanation, user protection, human management and oversight, and document preparation and retention. The EU AI Act regulates obligations for high-risk AI systems across seven articles (Articles 9 through 15).

Structure creates density. Combining Articles 9 through 15 of the EU AI Act, dozens of detailed requirements are listed directly within the statutory text. Korea writes only five broad duties into the statute, delegating fine details down to the Enforcement Decree, public notices, and guidelines. Under EU regulations, a business operator can know what needs to be done simply by reading the statute book. Under Korea's Basic Act, reading the statute book alone is insufficient; the complete picture emerges only after going through the Enforcement Decree, public notices, guidelines, and

consultations with the Support Desk.

Two obligations present in the EU are absent in Korea:

Data governance. Article 10 of the EU Act addresses quality criteria for datasets used in training, validation, and testing under a separate article. It lists eight items covering data collection processes and provenance, pre-processing operations (annotation, labeling, cleaning, augmentation), bias reviews, representation verification, and personal data protection exemptions. The Korean AI Basic Act contains no provision directly regulating the quality or governance of training data. While there is a provision requiring an explanation of an "overview" of training data, provisions establishing quality standards for data itself are missing.

Accuracy, robustness, and cybersecurity. Article 15 of the EU Act requires high-risk AI systems to achieve appropriate levels of accuracy, robustness, and cybersecurity. Resilience against errors and faults, defense against data poisoning and adversarial examples, and prevention of bias magnification through feedback loops are explicitly stated in statutory provisions. Korean law lacks a corresponding independent provision. Although it may be subsumed as part of risk management, there is a difference in binding force

between explicit statutory mandate and coverage at the level of public notices.

The assessment that this structure of Korean law is "flexible" is accurate. In a landscape where AI technology changes rapidly, fixing detailed standards in statute necessitates amending the law every time technology races ahead. Public notices and guidelines can be changed without statutory revision. However, there is a point where flexibility turns into uncertainty. If the law itself does not dictate what must be done, even business operators willing to comply do not know what to comply with. The 552 inquiries represent a snapshot of that uncertainty.

6. Between Recommendation and Obligation

The five duties under Article 34, Paragraph 1 carry the phrase "shall implement." They are obligations. Yet regarding specific matters prescribed by public notice of the Minister of Science and ICT, it reads that the Minister "may recommend compliance" (Article 34, Paragraph 2). They are recommendations.

A gap exists between obligation and recommendation. Implementing the five duties is a statutory obligation, but the detailed criteria on how to implement them are recommendations.

The expression "duty of autonomous effort" is used to describe this structure.

The administrative fine is up to 30 million won. Moreover, it is not imposed immediately upon violating a duty, but only when a business operator fails to comply with a suspension order or corrective order issued following a fact-finding investigation (Article 40, Paragraph 3; Article 43, Paragraph 1, Item 3). The grace period announced by the government is at least one year, extending beyond January 2027. Unless it involves grave matters such as casualties or large-scale human rights violations, administrative guidance and corrective recommendations precede administrative fines.

An amount of 30 million won is a burden to a 10-employee startup. To a large enterprise with annual revenue of 1 trillion won, it is 0.003 percent. There is a basis of comparison. Fines for violating high-risk AI obligations under the EU AI Act reach up to 15 million euros or 3% of global annual turnover. Converting 15 million euros into Korean won yields approximately 23 billion won—a nearly 800-fold difference from Korea's maximum administrative fine. The fine proposed in the Chinese scholars' draft is up to 50 million yuan (approximately 9 billion KRW). Korea's 30 million won is significantly lower among the

three countries.

The size of a fine does not automatically translate to regulatory effectiveness. Korean law incentivizes business operator behavior through other avenues.

Linkage with public procurement. Article 35, Paragraph 2 of the Basic Act dictates that state organs prioritizing the use of high-impact artificial intelligence products or services should preferentially consider those that have conducted impact assessments. Article 16, Paragraph 3, introduced via an amendment in January 2026, mandates state organs to preferentially consider artificial intelligence products and services when purchasing products or services required for their duties. This amended provision takes effect on July 21, 2026. Combining the two provisions yields the following structure: for business operators seeking government clients, carrying out impact assessments and fulfilling duties serves as a practical condition for market entry. What statute does not mandate, the market enforces. If the government procurement market is not large, this linkage is weak; as the government more aggressively promotes AI adoption, this linkage grows stronger.

Mechanisms to mitigate duplicate regulation carry practical significance for existing business operators. A business operator that has already implemented measures equivalent to each item of Article 34, Paragraph 1 pursuant to other statutes is recognized as having fulfilled the duties under the AI Basic Act in accordance with standards set forth in Attached Table 1 of the Enforcement Decree (Article 34, Paragraph 3). If a financial company operating a loan evaluation AI complies with the explanation duty under Article 35-2 of the Use and Protection of Credit Information Act and maintains explanation procedures, it is deemed to have fulfilled the explanation plan duty. If it fulfills all duties under Article 10 of the Financial Consumer Protection Act in the same domain, it is deemed to have fulfilled user protection plans as well. If a digital medical device manufacturer establishes a quality management system under the Digital Medical Products Act and its software receives a determination of compliance with quality management standards, it is recognized as having implemented risk management plans; if it satisfies all quality management system requirements under the *mutatis mutandis* applied Medical Devices Act, it is also recognized as having fulfilled documentation duties. In cases where measures under the Personal Information Protection Act have been implemented, the corresponding duties are deemed fulfilled

regarding personal information processing and protection.

It also reduces the burden on deploying business operators. If a developer business operator has fulfilled duties regarding high-impact artificial intelligence and the deploying business operator has not significantly altered the original purpose and intent of use, the deploying business operator is deemed to have fulfilled three duties: risk management plans, explanation plans, and user protection. Duties to operate human oversight systems and retain documents must be fulfilled by the deploying business operator itself. This division of responsibilities between developer business operators and deploying business operators stands in contrast to the EU approach, which distributes obligations across the entire supply chain.

The EU imposes conformity assessments on providers and operation and monitoring duties on deployers, while requiring fundamental rights impact assessments for public authorities and certain deployers such as credit scoring and insurance entities. Both sides bear substantial obligations. Korea concentrates primary responsibility on developer business operators while lowering the burden on deploying business operators. This design aims to relieve

the regulatory burden on adopting institutions (hospitals, schools, financial institutions). On the other hand, room remains for deploying business operators to reconfigure AI systems to suit their own contexts while arguing that they "did not significantly alter functionality." As this boundary remains blurred, it leaves room for future disputes.

On January 22, 2026, the five duties took legal effect and began applying to business operators. Public notices and guidelines were distributed concurrently. The grace period is underway. Mechanisms such as the right of deploying business operators to request materials from developer business operators, the duty of developer business operators to make cooperative efforts, and recognition of compliance pursuant to other statutes have begun to function. The structure created by law is running, but whether that structure is sufficiently dense is a matter for practice to judge. The questions piling up at the AI Basic Act Support Desk attest that this judgment is not yet complete.

Chapter 12. High-Impact AI Impact Assessment -- De Facto Essential, Not Legally Compulsory

In March 2026, an AI startup in Seoul received a bid announcement from an agency affiliated with the Ministry of Education. It was a project to build a middle school student grade prediction system. This is a high-impact AI system corresponding to the domain of "student evaluation in early childhood education, primary education, and secondary education" specified in Article 2, Subparagraph 4, Item (j) of the AI Basic Act. They open the bid documents. Following items like technical capability, security certification, and personal information processing track record, a single line catches their eye: "Execution of AI impact assessment." It is an extra point item. The law has never mandated impact assessments. Nor is there an administrative fine. Scoring zero on this item means falling behind competitors. The compliance team leader downloads the AI Impact Assessment Guidelines issued by the Ministry of Science and ICT.

1. Legal Nature of the Impact Assessment

A high-impact AI impact assessment is not mandatory.

Nowhere in the AI Basic Act is there a provision stating that an administrative fine will be imposed if an impact assessment is not conducted. Chapter 4 of the Act consolidates five duties into a single

article and imposes them as obligations on high-impact AI business operators—establishing risk management measures (Article 34, Paragraph 1, Subparagraph 1), preparing explanatory measures (Subparagraph 2), user protection (Subparagraph 3), human control and supervision (Subparagraph 4), and document preparation and retention (Subparagraph 5)—but the impact assessment itself is prescribed in Article 35, Paragraph 1 using the ending phrase "shall make efforts." In jurisprudence, such provisions are referred to as a "duty to endeavor." This means not that a result must strictly be achieved, but that efforts should be made toward that direction.

During the National Assembly's deliberation process, this provision originally took a different form. The initial bill included a provision imposing a prior impact assessment as a mandatory requirement. Business operators pushed back. They argued that requiring an impact assessment prior to releasing an AI system would delay market entry and cause them to fall behind in global competition. A compromise was reached. The structure leaves the impact assessment as a voluntary duty to endeavor, while granting preferential consideration in public procurement to business operators who carry it out. It selected the carrot of incentives instead of the whip of coercion.

In terms of legal textbook classification, this provision is close to a programmatic norm. It provides direction to the state and business operators, but no sanctions follow upon non-compliance. Even if an impact assessment is not conducted, there is no administrative fine and no administrative disposition.

Why, then, do it?

2. The Door Called Public Procurement

The answer lies in Article 16, Paragraph 3 and Article 35, Paragraph 2.

The amended Act (Act No. 21311), promulgated on January 20, 2026, newly established provisions to promote the adoption of AI in the public sector. This is Article 16, Paragraph 3. State agencies must give priority consideration to AI when purchasing products or services or placing service orders. Article 15, Paragraph 4 of the amended Enforcement Decree, which takes effect on July 21, 2026, defines the scope of priority consideration as AI products and services verified by the Minister of Science and ICT, delegating the standards and procedures for verification to a public notice.

Impact assessments fit into this structure. Article 35, Paragraph 2 of the Act stipulates that when state agencies intend to use products or services utilizing high-impact AI, they shall give preferential consideration to products and services for which an impact assessment has been conducted. Although the law does not compel it, undergoing the assessment is advantageous for entering the government market.

The annual volume of public procurement by the South Korean government exceeds 180 trillion won. Among this, the proportion of AI-related projects is rapidly expanding. As Article 16, Paragraph 3 of the amended Act codifies priority consideration for AI into law, the business-to-government (B2G) market becomes a non-negligible source of revenue for AI business operators. An immunity clause for civil servants was also included. Absent intent or gross negligence, civil servants bear no liability for damages arising from the adoption of AI. One reason for civil servants to hesitate to adopt AI has been removed. As barriers on the buyer side are lowered, the weight of the impact assessment as a qualification requirement on the supplier side grows accordingly.

Business operators who have not undergone an impact assessment are not barred from participating in bids. However, they are placed at a disadvantage compared to those who have. "You do not have to do it. You will simply be pushed back in the public market." This is the design of soft regulation.

3. Three Stages of the Impact Assessment

The AI Impact Assessment Guidelines released by the Ministry of Science and ICT on January 22, 2026, outline three stages. It is the last among the five major guidelines (Securing Transparency, Securing Safety, Determining High-Impact AI, Business Operator Obligations, and AI Impact Assessment).

First is the preliminary preparation stage. The purpose, function, and domain of application of the AI system subject to evaluation are finalized. It is determined whether the system qualifies as a high-impact AI. The determination criteria are the 11 domains under Article 2, Subparagraph 4 of the Act and the two-step procedure under Article 25 of the Enforcement Decree. In Step 1, domain applicability is verified: is it used in areas such as energy supply, healthcare, medical devices, nuclear facilities, biometric identification, recruitment and loan screening, means of transportation, public services, or student evaluation? In Step 2,

whether it exerts a significant impact on human life, bodily safety, or fundamental rights is reviewed. If the determination is difficult, a request for verification may be submitted to the Minister of Science and ICT, and the Minister must reply with the results within a default period of 30 days, which may be extended once by up to 30 days. An evaluation team is formed, and methods for gathering opinions from external stakeholders (users, affected individuals, civil society organizations) are established.

The main assessment stage follows. The core is the drafting of risk scenarios. It specifically depicts what happens when the AI system malfunctions or outputs biased results: a scenario where a loan screening AI systematically rejects residents of a specific region; a scenario where a student evaluation AI predicts grades for students with disabilities lower than their actual performance; a scenario where a recruitment AI scores female applicants' applications lower than male applicants'. Rather than abstract risks, one must specifically describe who suffers harm and how.

The evaluation items span multiple criteria: fairness (is it free from bias and discrimination?), transparency (can its operating principles be explained?), safety (are safeguard mechanisms in place in case of

malfunction?), accountability (is the locus of responsibility clear when problems occur?), personal information protection (is data processed lawfully?), and human controllability (can a human intervene to change the outcome?). The fundamental rights impact analysis examines impacts on constitutional rights such as the right to equality, the right to know, privacy, and the freedom of choice of occupation.

The post-assessment action stage concludes the process. Based on the evaluation results, measures to mitigate risks are formulated. A monitoring plan is established, and re-evaluations are conducted periodically. The evaluation results and action details are documented and retained. This meshes with the five-year retention obligation for compliance evidentiary documents required by Article 34, Paragraph 1, Subparagraph 5 of the Act and Article 27, Paragraph 2 of the Enforcement Decree.

4. Who Conducts the Assessment?

The business operator itself.

Self-assessment is the baseline. Article 28, Paragraph 2 of the Enforcement Decree stipulates that it may be conducted directly or commissioned to a third party, but there is no mandatory provision

requiring verification by an external certification body. Involving external experts is recommended by the guidelines, but it is not obligatory. The business operator writes the risk scenarios for its own system, scores the evaluation items itself, formulates post-assessment measures itself, and creates and retains the documents itself.

Two opposing views conflict over whether this structure is sound.

Those supporting self-assessment view it as follows: the party best suited to understand the internal workings of an AI system is the business operator who developed it. It is difficult for an external certification body to accurately grasp the bias of a complex machine learning model within a short audit period. Self-assessment enables practically deeper analysis. Mandating external certification increases costs and creates bottlenecks when there are insufficient specialized AI certification bodies in Korea.

The opposing side views it this way: if one assesses the risks of a system one created, an incentive arises to understate the risks. If the impact assessment yields poor results, it hinders business, so one might draft mild scenarios or fill in bias analyses as a mere formality. The law contains no external mechanism to check this.

Consider a scenario where a fintech company that developed a loan screening AI conducts a self-impact assessment. Suppose this company's model exhibits a pattern of statistically producing lower credit scores for residents outside the capital region. If it honestly discloses this pattern in its self-assessment, the risk as a high-impact AI is rated high, imposing additional burdens of explanation in transactions with public institutions. If it does not disclose it, nobody verifies it. This vulnerability in the incentive structure is not something intended by law, but a limitation inherent in self-regulation.

5. The EU's Conformity Assessment: Same Name, Different Edge

The conformity assessment under the European Union's AI Act is of a different nature.

Korea recommends it. The EU mandates it.

Article 43 of the EU AI Act mandates a conformity assessment for providers of high-risk AI systems prior to placing them on the market. Failing to pass the assessment means failing to receive CE marking, and without CE marking, entry into the EU market is impossible. The door is closed.

There are two pathways for conformity assessment. For most high-risk AI systems, the internal control method applies. The provider itself demonstrates compliance with the requirements of Articles 9 through 15 of the EU AI Act: risk management systems, data governance, technical documentation, record-keeping, provision of transparency information, human oversight mechanisms, and accuracy, robustness, and cybersecurity. All seven of these must be met. Even under internal control, technical documentation must be drafted, a quality management system must be operated, and an EU declaration of conformity must be issued.

A stricter pathway is set up for AI systems used in remote biometric identification. If harmonized standards have not been fully applied, an independent third-party certification body (notified body) directly audits the system. The business operator's self-assessment is no longer sufficient, and external eyes intervene.

Article 27 of the EU AI Act separately stipulates a Fundamental Rights Impact Assessment (FRIA). This is a distinct obligation from the conformity assessment. When public authorities or private entities providing public services deploy high-risk AI, they must conduct a fundamental rights impact assessment prior to deployment. This is

an obligation not of the provider, but of the deployer. It requires documenting the categories of affected persons, period of use, specific risks, human oversight measures, and mitigation plans for adverse impacts. The results must be notified to the national supervisory authority.

To summarize: the EU separated the two assessments—conformity assessment for providers, and fundamental rights impact assessment for deployers. Both are mandatory. Korea did not introduce an EU-style conformity assessment, but instead established a single impact assessment as a voluntary duty to endeavor.

6. Five Decisive Differences

Although Korea's impact assessment and the EU's conformity assessment share similar names, they differ in almost every aspect.

Their legal nature is different. In Korea, it is a duty to endeavor. In the EU, it is a legal prerequisite for market entry. If an impact assessment is not conducted in Korea, one suffers disadvantages in public procurement. If a conformity assessment is not conducted in the EU, one cannot enter the market.

The evaluating entity is different. In Korea, it is a self-assessment conducted directly by the business operator or commissioned to a third party, without mandating external verification. In the EU, while most are self-assessments (internal controls), third-party notified bodies may be brought in for sensitive areas such as remote biometric identification.

The format of the results is different. In Korea, the impact assessment results are internal documents retained by the business operator. Passing the EU's conformity assessment results in a physical CE marking affixed to the product and the issuance of an EU declaration of conformity—a visible signal that the market and consumers can see.

Sanctions for violations are different. In Korea, there is no administrative fine for failing to conduct an impact assessment. In the EU, placing high-risk AI on the market without a conformity assessment incurs fines of up to 15 million euros or 3 percent of global turnover. Violations of provider obligations, including improper affixing of CE marking, fall into the same tier. Violating prohibited AI practices escalates up to 35 million euros or 7 percent of turnover. Between the maximum administrative fine of 30 million

won across the entire AI Basic Act of Korea and the EU's 35 million euros, there lies a four-digit gap.

The timing of application is different. Korea's impact assessment guidelines have been in effect since January 22, 2026. The EU's high-risk AI conformity assessment obligation is expected to be delayed until December 2, 2027, once the Digital Omnibus amendment is finalized. While Korea moved first, once the EU actually begins enforcement, the weight of regulation will be overwhelmingly heavier on the EU side.

7. Why Korea Opted for Self-Regulation

It was an intentional choice.

The official title of the AI Basic Act is the "Basic Act on the Development of Artificial Intelligence and Creation of a Foundation for Trust, etc." "Development" comes before "Trust." This law is a hybrid of promotion and regulation, with the center of gravity placed on promotion.

During National Assembly deliberations, the industry's logic was as follows: the Korean AI industry is a latecomer in the global market. Imposing prior conformity assessments on domestic startups and

SMEs that must compete with large AI enterprises from the United States and China would slow market entry speeds and increase costs. Even the EU proposed delaying high-risk AI obligations by 16 months until December 2027 under the Digital Omnibus amendment, aiming for a 35 percent reduction in regulatory compliance costs for SMEs. The argument went: while the EU itself slows down its regulatory pace, why should Korea set up strong regulations ahead of them?

The Ministry of Science and ICT took the same direction. In forming a subordinate legislation task force of around 80 private-sector experts, it established the business operator's voluntary securing of safety as a foundational principle. Five guidelines and two public notices (Public Notice on Securing Safety of Frontier AI, Public Notice on Obligations of High-Impact AI Business Operators) adhere to this principle. Guidelines carry no legal binding force. They serve as practical manuals demonstrating how business operators can fulfill their statutory obligations.

The level of administrative fines shares the same context. A maximum of 30 million won. Underlying this is the assessment that the domestic AI industry is in an early growth stage. Concerns that high fines could stifle innovation were continuously raised

throughout legislative deliberations.

The incentive structure stems from this philosophy as well. Instead of punishing through law, the intention is to draw operators in through the market. Performing an impact assessment yields advantages in public procurement. Anonymous consulting can be obtained at the AI Basic Act Support Desk. During the grace period (at least one year, through January 2027), administrative guidance and corrective recommendations take precedence over administrative fines. It is a structure that lays out multiple carrots while temporarily setting aside the whip.

8. Flaws in This Design

Criticism arises from multiple quarters.

In a September 2025 report ("Study on Rational Regulation and Development Plans for the AI Basic Act"), Senior Research Fellow Won-jun Jung of the Korea Legislation Research Institute pointed out that the concept of "high-impact AI" lacks precision. Within the Act, "impact" and "risk" are used side by side, while "safety" and "fundamental rights" are linked by "and," causing interpretive confusion. Even if a business operator attempts to perform an impact assessment, if the criteria for what constitutes a "significant impact"

remain ambiguous, the assessment itself devolves into a formality.

The same report highlighted the problem that the scope of regulated entities is overly broad, encompassing even government agencies within "AI business operators." This stems from failing to subdivide duty-bearers according to the value chain (provider, distributor, deployer, user) as the EU AI Act does, leaving only two broad categories: developer business operators and user business operators. The subject and scope of impact assessments rest upon this ambiguity.

In a report published on the date of the Act's enforcement, January 22, 2026 ("Administrative Tasks for Realizing the National AI Strategy"), Legislative Research Officer Jun-hwa Jung of the National Assembly Research Service stated that specifying the concept and requirements of high-impact AI and differentiating obligations by business operator type remain pending tasks.

In an October 2024 report ("Legislative Study for AI Impact Assessment"), Research Fellow Yoo-bong Lee of the Korea Legislation Research Institute comparative-analyzed the impact assessment systems of the EU and the US, recommending that to ensure the effectiveness of the Korean-style impact assessment, the scope of

disclosure for evaluation results and the external verification framework must be designed in detail.

The self-assessment structure itself harbors vulnerabilities. Even if a business operator intentionally downplays risk scenarios or fills in bias analyses format-wise, the law lacks any mechanism to verify it. This stands in contrast to the EU requiring third-party certification body evaluations for remote biometric identification AI.

There is also the deterrence problem of the 30 million won administrative fine. For a global AI enterprise generating hundreds of billions of won in annual revenue, 30 million won is not even perceived as a cost of doing business. The EU setting the cap at 7 percent of turnover stems from the judgment that sanctions must be heavy enough to affect business decisions for regulation to work.

9. Does Soft Regulation Work?

The answer to this question has not yet emerged.

Prerequisites exist for soft regulation to function: market participants must understand the intent of the law and follow it voluntarily, or the market disadvantages of non-compliance must be sufficiently large. The public procurement market must be large

enough, and whether an impact assessment was conducted must carry substantive weight in procurement evaluations.

Proponents argue as follows: the Korean government's AI policy is moving on all fronts. The National AI Strategy Commission finalized an action plan containing 99 execution tasks and 326 policy recommendations. Vice-minister-level Chief AI Officers (CAIOs) have been designated across all 48 central administrative agencies, holding monthly regular council meetings. An AI Policy Center was established within the NIA, with 59 personnel performing policy development and impact analysis. Once this framework is operational, whether an impact assessment has been executed naturally becomes an important evaluation criterion in AI procurements by public institutions. Even if the law does not compel it, the market will.

Opponents respond as follows: regulation relying on market autonomy fails in areas where the market does not function. For business operators uninterested in the B2G market or operating exclusively business-to-consumer (B2C) services, public procurement incentives have no effect. For large conglomerates running recruitment AI internally or financial institutions using loan screening

AI strictly in-house, public procurement extra points exert zero influence. Motivation for these operators to voluntarily conduct impact assessments does not exist within the law.

The track record of the AI Basic Act Support Desk provides reference. Within 10 weeks of enforcement, 552 consultation requests were received. Nearly half of the inquiries concentrated on transparency obligations, followed by high-impact AI verification. This signals an intention among business operators to comply with the law. However, consultation counts and actual impact assessment execution counts are separate matters. There is a gap between asking and executing.

Comparing the paths of the three jurisdictions, each choice is distinct. The EU chose hard regulation: without passing a conformity assessment, market entry is impossible, and violations trigger fines calculated as a percentage of global turnover. China chose prior registration (备案): releasing generative AI services with public-opinion shaping power or social mobilization capacity requires registering the algorithm within 10 business days. As of November 2025, 611 services completed registration. Failure to register can lead to demands for rectification and ultimately service

suspension. Korea chose incentives: non-compliance incurs no punishment, while compliance offers advantages in the public market.

Which approach is correct cannot be judged in July 2026. The EU's high-risk AI conformity assessment obligations take full effect only in December 2027. Korea's public procurement linkage structure takes effect on July 21, 2026. Empirical comparison will only be possible years down the road.

Yet one trajectory is legible: as CAIOs from 48 agencies meet monthly, 99 tasks are pursued, and Article 16, Paragraph 3 codifies priority consideration for AI into law, the likelihood of an AI system that has not undergone an impact assessment being selected in public tenders will shrink with each passing year.

That startup in Seoul begins its impact assessment. Not because the law demands it, but to win the bid. The competitor next door is making the same calculation. This is the moment a regulation without coercion goes to work. Whether that mechanism can expand into markets beyond B2G, whether it can prevent self-assessments from becoming mere formalities, and whether the 30 million won administrative fine can possess sufficient deterrence—these

questions left open by the law await answers in the market.

Chapter 13. Procedure for Requesting Confirmation of High-Impact Status

In March 2026, in a meeting room of a fintech company in Gangnam, Seoul, the legal team lead projected a slide. "Our loan screening AI is indeed high-impact." The development team lead retorted, "Since it's in the loan screening domain, isn't it automatically high-impact? Why do we need to make a separate request for confirmation?" The legal team lead shook their head. "Falling under the domain does not immediately make it high-impact. There is a two-stage determination. And if we don't request confirmation, we will get into trouble later." The meeting room fell silent. Half of the people in the room were learning for the first time on the spot that they had to submit a confirmation request form to the Ministry of Science and ICT.

1. What Is a Confirmation Request?

The AI Basic Act imposes responsibilities on business operators operating high-impact AI, such as establishing risk management plans, the duty of explanation, securing human intervention systems, and retaining documents for five years. Whether a business operator must assume these responsibilities depends on whether their service falls under high-impact status. A confirmation request is the official

procedure for inquiring with the Ministry of Science and ICT as to whether it falls under such status.

The business operator is the primary actor. The structure is not one where the government approaches first and notifies, "Your AI is high-impact." Article 33, Paragraph 1 of the Act imposes an obligation on business operators to review in advance whether their artificial intelligence constitutes high-impact AI. When they remain uncertain even after conducting the review, they request confirmation from the Minister of Science and ICT. In this respect, the South Korean system uses the business operator's prior review as its starting point while providing a supplementary mechanism where the government provides an answer when determination is difficult.

The targets eligible to request confirmation include both development business operators and utilization business operators. Development business operators who wonder whether their developed AI models are being used in high-impact domains, as well as utilization business operators who receive external models and apply them to loan screenings or hiring evaluations, can undergo the exact same procedure.

2. Flow of the Procedure: 30 Days, and Another 30 Days

Once a confirmation request is received, the Ministry of Science and ICT notifies the results within a basic period of 30 days. Taking into consideration factors such as the complexity and importance of the product or service, this period may be extended by 30 days on a one-time basis, in which case the reasons for the extension and the extension period must be notified in writing. The maximum processing period is 60 days.

The 30-day period was discussed multiple times during the enactment process of the Enforcement Decree. Industry stakeholders argued that it should be reduced to 15 days. Their rationale was clear: in an environment where AI service release cycles are accelerating, 30 days is close to a whole generation. If a business operator receives confirmation results after a competitor has already launched their service, market opportunities are lost. For startups, a 30-day gap could collide with fundraising schedules.

There were counterarguments as well. Determining high-impact status requires examining technical characteristics and social ripple effects simultaneously. AI in the medical domain and AI in the energy supply domain have completely different safety standards that must be reviewed. Expert consultation may also be required. The

counterargument maintained that it is difficult to precisely evaluate all domains within 15 days.

As a result, the basic 30 days was maintained. However, as there is no lower limit on the deadline under the law, clear-cut cases with few contentious issues may receive an answer in 10 days without taking the full 30 days.

3. Structure of the Two-Stage Determination

The determination criteria applied by the Ministry of Science and ICT upon receiving a confirmation request consist of two stages.

Stage 1 is domain verification. It examines whether the relevant AI system is used in one of the 11 domains enumerated in Article 2, Item 4 of the Act: energy supply; drinking water production processes; healthcare; medical devices and digital medical devices; safety management of nuclear materials and nuclear facilities; analysis and utilization of biometric information for crime investigation or arrest operations; decisions significantly affecting individual rights and duties such as recruitment or loan screening; management of major operations of transportation means and transportation facilities; decision-making by state agencies, etc. affecting citizens such as verification and determination of qualifications required for public

service provision and cost collection; student evaluation in early childhood education, primary education, and secondary education; and domains that may be additionally designated by Presidential Decree.

Stage 2 is the determination of materiality. The mere fact that it is used in a specified domain does not automatically make it high-impact. There must be a "concern of significantly affecting or causing danger to human life, physical safety, and fundamental rights." Even if AI is used in the loan screening domain, if that AI does not make the final decision but merely serves as an auxiliary tool organizing reference material, it may not satisfy the materiality criterion. Conversely, if the AI effectively determines loan approvals and rejections and human personnel rarely overturn the results, materiality is highly likely to be recognized.

The Guidelines for Determining High-Impact AI present specific criteria that business operators can reference for this two-stage determination. Although the guidelines themselves lack legal binding force, because the Ministry of Science and ICT uses them as the basis for judging confirmation requests, they serve as a de facto benchmark norm.

4. Discussions on Self-Checklists and Low-Impact Notification Systems

During the legislative notice process of the Enforcement Decree, an alternative proposal emerged as feedback: a discussion on introducing a simplified notification system for AI that does not constitute high-impact AI—that is, low-impact AI.

The concept is as follows: The Ministry of Science and ICT provides a self-check list. The business operator checks this list item by item, and if they make an internal determination that their AI does not belong to any of the 11 domains, or even if it does, it fails to meet the materiality criteria, they notify the Ministry of Science and ICT of the result. Rather than going through the 30-day confirmation procedure, this approach allows confirmation of 'not high-impact' simply through notification.

Behind this discussion lay practical concerns. The number of business operators providing AI services in South Korea is surging rapidly. If every business operator submits a confirmation request to the Ministry of Science and ICT, administrative capacity will be overwhelmed. Subjecting even clearly low-impact matters to a 30-day procedure represents inefficiency for both business operators and the government.

As of July 2026, self-check lists are partially reflected in the guidelines, but an official low-impact notification system has not been institutionalized. If a business operator independently determines that its AI is not high-impact, retaining the grounds for that judgment in writing is a practically safe response. This is because, should a fact-finding investigation occur later, the business operator must prove why it deemed the AI not high-impact.

5. Legal Effect of Confirmation Results

When the Ministry of Science and ICT responds that the AI "constitutes high-impact AI," the business operator must fulfill the obligations under each item of Article 34, Paragraph 1 of the Act. They must establish and operate risk management plans. They must formulate and implement plans to explain the final results derived by the AI, the major criteria used to derive those results, and an outline of the training data. They must establish and operate user protection measures. They must secure a system for human management and supervision. They must prepare documents verifying the details of measures taken and retain them for five years. In addition, pursuant to Article 27, Paragraph 1 of the Enforcement Decree, they must post the main details of the risk management plans, the main details of the explanation plans, user protection measures, and the name and

contact information of the person managing and supervising on their office, place of business, or internet homepage, with matters constituting trade secrets eligible for exclusion.

What legal effect does a response stating "does not constitute high-impact AI" hold? It is a confirmation that, based on the specific service format at that point in time, the business operator is not required to bear obligations for high-impact AI. It is not an exemption. If the service changes, the determination may change as well. If an AI that served as an auxiliary tool for loan screening transforms into an effectively automated screening system through a version upgrade, the business operator must re-examine whether it constitutes high-impact AI based on the modified service, and it is safe to request confirmation again if necessary.

When a business operator disagrees with the determination that its AI is high-impact, it may first request a re-confirmation. Pursuant to Article 25, Paragraphs 4 and 5 of the Enforcement Decree, if a re-confirmation request form is submitted within 10 days from the date of receiving the response, the Minister of Science and ICT consults an expert committee and responds with the results within 30 days from the date of receiving the re-confirmation request. If the

business operator still disagrees with the re-confirmation results, whether administrative litigation is possible becomes an issue. The general interpretation among legal experts is that if the confirmation response constitutes an administrative disposition, it can be contested through administrative appeals or administrative lawsuits.

6. Subsequent Procedures After Being Confirmed as High-Impact

If the confirmation result is "high-impact," the business operator must sequentially undertake several measures.

Establishing a risk management plan is the first step. The Ministry of Science and ICT's guidelines present the risk management process in five stages: policy formulation, risk identification, analysis and evaluation, treatment, and improvement. This is similar to the structure of ISO 31000 (the international standard for risk management). It is a cyclical structure of identifying what risks one's AI system inherently entails, analyzing the severity and probability of occurrence for each risk, executing measures to mitigate risk, and monitoring and improving the outcomes.

Impact assessment is recommended. For high-impact AI business operators, fundamental rights impact assessment is not a mandatory obligation but an 'obligation of voluntary effort.' Although not legally

enforced, business operators that perform impact assessments receive priority consideration in public procurement. For business operators seeking to enter the government-facing market, impact assessment is effectively essential.

Transparency obligations accompany this. The fact that high-impact AI is being operated must be disclosed to users, and an explanation plan must be established and implemented so that affected individuals can receive explanations regarding the main criteria and principles behind the AI's determinations.

Documentation obligations have a long temporal tail. Evidentiary documents of implementation must be retained for five years. Article 27, Paragraph 2 of the Enforcement Decree does not specify a separate starting point for the retention period; therefore, even if a service is terminated, it is safe to retain the documents until five years have passed from the time each obligation was fulfilled.

7. EU Conformity Assessment: The Dual Structure of Self-Declaration and Third-Party Assessment

The European Union AI Act's high-risk AI regime has a point of fundamental divergence from South Korea. The EU does not make business operators ask the government to "please confirm whether

your AI is high-risk." The business operator must judge for itself and pass a conformity assessment before placing the AI on the market.

There are two pathways for conformity assessment. One is self-declaration. The business operator internally declares that it satisfies the requirements of the EU AI Act, prepares technical documentation, and affixes the CE marking. Most high-risk AI systems listed in Annex III follow this pathway.

The other is third-party conformity assessment. An independent conformity assessment body (Notified Body) inspects the relevant AI system. The scope where this pathway is required is narrow. Cases among biometric systems in Annex III that have not fully applied harmonized standards or common specifications, and types linked with existing product safety legislation such as machinery or medical devices, fall under this category. Other high-risk areas such as critical infrastructure or law enforcement follow the internal control method—namely, the self-declaration pathway—in principle.

South Korea does not fall into either of these two pathways. South Korea's high-impact confirmation is a structure where the business operator poses a question to the government, whereas the EU's conformity assessment is a structure where the business operator

directly proves compliance or a third party verifies it. The philosophical difference—where South Korea focuses on impact and the EU focuses on risk—translates here into a procedural difference.

The disparity in the scale of sanctions underpins this structural difference. In the EU, placing high-risk AI on the market without a conformity assessment can incur fines up to 15 million EUR or 3% of total worldwide annual turnover, whichever is higher. South Korea is different. Administrative fines under Article 43 of the Act are imposed only when prior notification is not given, a domestic representative is not designated, or a suspension order or corrective order is not fulfilled, with a statutory cap of 30 million KRW. Rather than administrative fines attaching directly to a violation of high-impact business operator duties per se, the structure leads to administrative fines only when a corrective order is issued and remains unfulfilled. This represents the gap between a regime requiring proof prior to market entry and a regime permitting confirmation post-market entry.

8. 17 Percent at the Support Desk

The AI Basic Act Support Desk has been taking questions from business operators since its opening on January 22, 2026. The consultation history reveals the reality of the law.

Among the 94 online inquiries received during the first ten days after opening, the most frequent category concerned transparency obligations, accounting for 56.4 percent. Questions included: "In what format should we disclose that we are using AI?" and "Where and how should we attach generative output labeling?" The second most frequent category was high-impact AI confirmation at 17 percent, and the third was consultation regarding the definition of business operators at 10.6 percent, consisting of questions like "Are we a development business operator or a utilization business operator?"

The structure remains identical across the 552 cumulative consultations recorded up to the 10th week of enforcement. Transparency-related inquiries still accounted for about half at approximately 51 percent, while high-impact confirmation increased in proportion to approximately 19.6 percent, representing about 108 cases.

What these figures—progressing from 17 percent to 19.6 percent—tell us is meaningful. It means that nearly two out of every ten business operators are unable to determine whether their service falls under high-impact status. Most of these business operators have

likely read the list of 11 domains under Article 2, Item 4 of the Act. The reason they cannot reach a determination even after reading the list is the Stage 2 materiality determination. While their service falls within a domain, it is difficult for them to independently settle whether their AI exerts a "significant impact" on fundamental rights.

The Support Desk operates on weekdays from 9:00 AM to 6:00 PM. Matters requiring in-depth legal review take up to 14 days, or up to 21 days if review is delayed. Consultations can also be conducted anonymously using pseudonyms without disclosing real names or company names. It provides a channel through which business operators can verify the legal status of their service without revealing trade secrets.

9. Risks for Business Operators Delaying Confirmation Requests

Requesting confirmation is not mandatory. However, because Article 33, Paragraph 1 of the Act establishes the act of reviewing in advance whether an AI constitutes high-impact as an obligation of the business operator, bypassing even the review is not what the law intends. If a business operator conducts its own review, determines that "our service is not high-impact," and refrains from submitting a confirmation request, that in itself is not unlawful. The problem arises when that determination turns out to be wrong.

What happens if, after a business operator makes an internal determination that its AI is not high-impact, a subsequent fact-finding investigation by the Ministry of Science and ICT reveals that it does constitute high-impact AI? That business operator enters a state of non-compliance with all obligations imposed on high-impact AI business operators. There is no risk management plan, no explanation system, and no retained documentation. While administrative guidance and corrective recommendations are likely to be prioritized during the early implementation period, if a corrective order is issued pursuant to Article 40 of the Act and remains unexecuted, it can lead to an administrative fine of up to 30 million KRW.

A greater risk lies in business costs rather than legal sanctions. Retroactively building compliance with high-impact AI business operator obligations costs several times more than establishing it from the outset. A new risk management plan must be formulated, documentation for past periods may prove impossible, and service designs may need to be altered. Establishing a compliance framework while simultaneously running a service is akin to changing tires on a moving car.

Disadvantages also emerge in the public procurement market. Article 35, Paragraph 2 of the Act provides that state agencies, etc. shall give priority consideration to products and services that have undergone impact assessments; thus, business operators that have failed to even resolve their high-impact status and consequently left impact assessments untouched risk falling behind in competition.

There is yet another risk. When users or affected individuals file civil claims for damages, the fact that a business operator did not request high-impact confirmation can be used as evidence of negligence. The underlying logic is that "a reasonable business operator would have verified whether its service constitutes high-impact AI." While there is no legal precedent yet on how courts will rule on this, failing to request confirmation is unlikely to serve as a favorable circumstance for the business operator.

10. Practical Application of the Guidelines for Determining High-Impact AI

The guidelines were released on January 22, 2026, as one of five major guidelines. They lack legal binding force. They are neither administrative rules nor public notifications. However, if the Ministry of Science and ICT examines confirmation requests, conducts fact-finding investigations, and makes determinations in accordance with the guidelines, business operators cannot afford to ignore them.

The decision-making flow presented by the guidelines operationalizes the two-stage structure of the Enforcement Decree. A business operator first specifies the intended purpose and application domain of its AI system. It checks whether any of the 11 domains apply. If a relevant domain exists, it proceeds to the next stage. It evaluates what role the AI system plays in that domain, to what degree it participates in final decisions, and at what level human oversight operates.

Grounds for determination must be recorded in writing. The guidelines recommend recording the grounds for "why it was determined to be high-impact" or "why it was determined not to be high-impact." This document serves as explanatory material for confirmation requests later on, as well as evidence demonstrating whether the business operator's self-determination was reasonable during fact-finding investigations.

A zone of tricky judgment in practice occurs when AI is used as an auxiliary tool yet deeply participates in decision-making in effect. Suppose there is an AI that scores and ranks resumes during a hiring process. The final hiring decision is made by an HR manager. However, if the HR manager has almost never overturned the AI's

rankings, this AI is effectively making the decision. The guidelines advise making determinations based on such substantive functions.

11. Structural Differences Between South Korean and EU Procedures

Placing the two regimes side by side reveals how differences in regulatory philosophy translate into procedural differences.

South Korean business operators ask the Ministry of Science and ICT, "Is this high-impact?" The government answers. EU business operators judge for themselves, undergo a conformity assessment, and declare to the market, "This is safe." South Korea makes the government the arbiter, whereas the EU makes the business operator the prover.

In the EU's self-declaration pathway, the business operator prepares technical documentation, establishes a quality management system, and signs the EU Declaration of Conformity. Only when these documents are fully prepared can the CE marking be affixed and the product placed on the market. Market surveillance authorities inspect these documents ex post facto.

In the EU's third-party assessment pathway, a Notified Body directly inspects the AI system. Types with high risk of fundamental rights

infringement, such as biometric systems, undergo this pathway. This is underpinned by the premise that a business operator's self-declaration alone is insufficient.

Neither of these pathways exists in South Korea. The very concept of conformity assessment is not introduced in the AI Basic Act. South Korea's confirmation procedure is limited to distinguishing "whether it is high-impact or not," without extending to proving "whether it is safe or unsafe." Instead, South Korea seeks to ensure safety through legal obligations under Article 34, Paragraph 1 of the Act, including the establishment of risk management plans, alongside the voluntary duty of effort known as impact assessment.

It is premature to conclude which approach is superior. The EU approach filters safety prior to market entry, but imposes a heavy burden on business operators and may slow the pace of innovation. The deferral of compliance deadlines under the Omnibus amendment—to December 2027 for standalone high-risk AI and August 2028 for AI embedded in products—serves as evidence of that burden. The South Korean approach does not block market entry, but if a problem erupts post-entry and the business operator's autonomous determination was flawed, enforcement sanctions are

weak. The question of whether the 30 million KRW administrative fine cap can exert a deterrent effect on large enterprises remains unresolved.

12. When Determination Is Uncertain

The core of requesting confirmation of high-impact status lies in the "grey zone where determination is uncertain." In obvious cases, such as AI for safety management of nuclear facilities or AI for biometric recognition in crime investigation, the confirmation request itself may not be necessary. The business operator can independently determine it to be high-impact and fulfill its obligations.

The problem lies in the grey zone. Is an AI tracking student learning progress in the education sector high-impact? Article 2, Item 4, Subitem (j) of the Act enumerates "student evaluation in early childhood education, primary education, and secondary education." Is learning progress tracking considered evaluation, or monitoring? What about an AI recommending diets in the healthcare domain? It could fall under the "establishment and operation of healthcare delivery and utilization systems," but can diet recommendations be viewed as significantly affecting human life or physical safety?

Such borderline cases account for approximately one-fifth of consultations at the Support Desk. As of July 2026, six months after the enforcement of the Act, judgment criteria will become more concrete as cases undergoing the confirmation request procedure and their outcomes accumulate. In the early stages, actual precedents of confirmation results serve as more useful reference material for business operators than the wording of guidelines.

A confirmation request is an insurance policy, not a cost. A business operator that has secured an official response from the Ministry of Science and ICT possesses grounding in situations where it must later prove the legitimacy of its determination. Business operators that delay requesting confirmation have no choice but to hope that their own judgment was correct.

Part 5. Adjacent Law Risks Encountered by Business Operators

Chapter 14. Training Data and Copyright: Fair Use Provisions, First Domestic Lawsuits, and Warnings from Overseas Precedents

On December 27, 2023, The New York Times filed a complaint in the U.S. District Court for the Southern District of New York. OpenAI and Microsoft were listed in the defendants column. Among the nearly 100 pages of evidence attached to the complaint, what caught attention was a screenshot showing ChatGPT reproducing a New York Times article almost word for word. An investigative report covering the taxi medallion crisis matched down to sentence units. Counsel listed the damages amount as "billions of dollars." It was the moment when the copyright issue of data used for AI training arrived in the courtroom.

1. A Single Line Called Training Data

Article 2, Subparagraph 12 of the AI Basic Act defines "training data." "It refers to data used for the development and utilization of artificial intelligence." It is a single line. Among the twelve definitions, it has the fewest characters and leaves the broadest blank area.

This definition does not limit the types of data. Text, images, audio, video, code, medical records, legal documents, news articles, novels.

If it is data put into creating or using artificial intelligence, it is all training data. Whose it is, how it was collected, whether permission was granted—this definition does not ask.

Article 15, Paragraph 4 of the AI Basic Act established the legal basis for constructing an integrated provision system for training data, and Article 13 of the Enforcement Decree sets the functions and management method of this system. It must possess functions for integrated search, systematic classification, and provision of sources for training data. It is an infrastructure providing data to be used for AI training, centered around public data. However, the AI Basic Act does not state what conditions are necessary when using non-public, private works for training. There is no link between the definition of "training data" under the AI Basic Act and the protection framework under the Copyright Act. Lawsuits have begun in this empty space.

2. Four Scales Called Fair Use

Article 35-5 of the Copyright Act provides for the "fair use of works." This is a provision introduced in the 2011 amendment to the Korean Copyright Act, adopting the fair use doctrine of Section 107 of the U.S. Copyright Act.

It is Paragraph 1. "A work may be used if it does not conflict with the normal exploitation of the work and does not unreasonably prejudice the legitimate interests of the author." Paragraph 2 requires four factors to be considered when determining this: The purpose and character of the use; The nature and purpose of the work; The amount and substantiality of the portion used in relation to the work as a whole; The effect of the use of the work upon its potential market or value of the work.

Applying these four criteria to AI training does not yield a clear answer.

Purpose and character of use. AI model training is commercial. OpenAI, Naver, and Kakao all collect data for commercial services. For-profit use is disadvantageous to establishing fair use. On the other hand, there is a counterargument that the training process extracts statistical patterns rather than enjoying the expression of the work itself. The argument is that if it is "analyzing" rather than "reading" the work, the character of the use is transformative.

Nature and purpose of the work. News articles, novels, academic papers, and photographs each have different degrees of protection. Articles close to factual reporting have a narrow scope of protection,

while novels or poems have a broad scope of protection. Because AI training collects large quantities regardless of genre, this criterion is not applied uniformly.

Amount used. AI models replicate whole works. When feeding a newspaper article into training, they input the full text, not a part of it. Entire replication is disadvantageous to establishing fair use.

Market impact. The fourth criterion is the sharp edge. If AI generates and provides similar content to users after learning news articles, the user's reason to visit the original news site diminishes. It is the logic that the potential market for the work is eroded. AI business operators counter that there is no direct substitution relationship between training data and generated outputs.

Korean courts have not yet rendered a ruling on whether AI training data constitutes fair use. Where the indicators on the four scales will be set will be determined when a ruling comes out.

There is something to note. The Korean Copyright Act does not have a provision that explicitly permits AI training or text and data mining (TDM), unlike Article 30-4 (Use for Purpose of Data Analysis) of the Japanese Copyright Act. There is only a general fair use provision

(Article 35-5). Because a general provision requires courts to weigh the four criteria on a case-by-case basis, it fails to provide predictability to business operators.

3. The Battlefront of the Seoul Courts

In Korea, broadcasting companies fired the first shot in copyright disputes over AI training data.

In January 2025, three terrestrial broadcasting companies—KBS, MBC, and SBS—filed a copyright infringement lawsuit against Naver in the Seoul Central District Court. The broadcasters' claim is that Naver used large quantities of broadcast news content without permission while training Korean large language models such as HyperCLOVA X.

In February 2026, the same three broadcasters filed a lawsuit against OpenAI in the Seoul Central District Court. The cause of action is that ChatGPT used broadcast news content as training data without obtaining permission. The battlefront has expanded from a domestic portal to an overseas AI company.

The questions this dispute poses to copyright law are clear. Can fair use under Article 35-5 apply to commercial AI training? Does AI

generating similar news summaries or answers after using large quantities of news content for training replace the market for the original work?

Portal business operators can argue that AI training constitutes fair use because it is an act of extracting facts and patterns rather than expressions of works. Broadcasters will counter that market erosion occurred because their content was used as core raw material for AI services, and those AI services replace news consumption.

The result of this lawsuit will decide the data procurement structure of the Korean AI industry. If the court recognizes fair use, content disclosed on the Internet can be used relatively freely for training. If fair use is denied, individual licenses must be secured for all training data. That cost is a burden even for large corporations and a wall for startups.

4. Two Wars in New York

In the United States, two lawsuits are at the center of the debate on AI training data copyright.

The New York Times v. OpenAI. Filed in the U.S. District Court for the Southern District of New York on December 27, 2023. The New York

Times claimed that OpenAI and Microsoft used millions of articles without permission to train GPT models. The core evidence in the complaint consisted of instances where ChatGPT output New York Times articles almost verbatim. AI even reproduced articles behind paywalls. The New York Times demanded billions of dollars in damages and a permanent injunction deleting its data from GPT models.

OpenAI put forward a fair use defense. Their logic is that even if full articles are reproduced, it is transformative use because the training process learns the statistical structure of language rather than the expression of the work. They countered that instances of article reproduction were model malfunctions (regurgitation), not designed functions.

While the lawsuit was ongoing, OpenAI began entering into licensing agreements with other news organizations. The Associated Press, Axel Springer (Germany), and Le Monde (France) executed content usage agreements. News organizations that sue and news organizations that contract coexist. It means that training data licensing is becoming an unavoidable cost for AI business operators.

Authors Guild v. OpenAI. In September 2023, the Authors Guild and famous authors such as John Grisham, Jodi Picoult, and George R.R. Martin filed a class action lawsuit in the U.S. District Court for the Southern District of New York. The cause of action was that GPT models used thousands of books for training without permission. Their argument is that their books were copied in full and put into AI, which constitutes an infringement of the reproduction right.

What is at issue in this lawsuit is the doctrine of transformative use. In the 2023 case *Andy Warhol Foundation v. Goldsmith*, the U.S. Supreme Court narrowed the scope of transformative use. It became harder to claim transformative use when using original works for commercial purposes. Although the Warhol ruling was not a case judging AI training itself, it can be cited as an argument to narrow the transformative use defense against large-scale reproduction for commercial purposes.

However, two rulings from the U.S. District Court for the Northern District of California in 2025 showed a different direction. In *Bartz v. Anthropic*, the court ruled that training language models with lawfully purchased books itself is a highly transformative fair use. However, downloading and storing pirated books was seen as a

separate infringement, and in September 2025, Anthropic agreed to a settlement of approximately \$1.5 billion with the authors. In *Kadrey v. Meta*, the court also accepted Meta's fair use defense in relation to the plaintiffs concerned. The trend is to recognize transformativeness at the training stage, while separately examining how data was acquired and what impact it has on the market.

As of July 2026, neither of the two New York lawsuits has reached a final judgment. However, the lawsuits themselves are already changing the market. Global AI companies have begun reflecting training data licensing costs into their business plans, and some companies are adopting strategies to transition to synthetic data.

5. The Door Kept Open by Tokyo

Japan maintained a legal framework more generous toward the issue of AI training data than any other country.

Article 30-4 of the Copyright Act of Japan provides that a work may be used in cases "where the use is not for the purpose of enjoying or letting others enjoy the thoughts or sentiments expressed in the work." It was introduced in the 2018 amendment of the Copyright Act. Data analysis, machine learning, and text and data mining (TDM) have been cited as typical application cases of this provision.

The logic goes as follows: when AI learns a novel, it does not enjoy the expression of the novel. It extracts statistical patterns of language. Because enjoyment is not the purpose, it falls under a limitation on copyright. According to this interpretation, large-scale reproduction for AI training is permitted in principle.

In 2024, the Copyright Subcommittee of the Council for Cultural Affairs under the Agency for Cultural Affairs of Japan re-examined the relationship between AI and copyright. There were two key issues: Does Article 30-4 apply even when AI generates outputs similar to the works it learned? Does it fall under cases that "unreasonably prejudice the interests" of the copyright owner (the proviso to the same Article)?

The Agency for Cultural Affairs released a report titled "Considerations Regarding AI and Copyright." It proposed a dual structure: allowing reproduction at the training stage under Article 30-4, while outputting expressions similar to original works at the generation stage could constitute copyright infringement. The approach is to leave training open while filtering generation.

There is a lesson that Japan gives to Korea. Japan established an explicit provision, yet fell into debate again over its scope. Korea has

no explicit provision itself. Handling AI training solely with a general fair use provision entails greater legal uncertainty than Japan.

6. Brussels' Principle of Disclosure

The European Union chose a direction demanding transparency on the issue of training data.

Article 53, Paragraph 1 of the EU AI Act imposes an obligation on providers of general-purpose AI (GPAI) models to draw up and implement a policy to comply with EU copyright law. Specifically, they must identify and respect the reservation of rights by copyright holders pursuant to Article 4, Paragraph 3 of the Digital Single Market Directive (DSM Directive, 2019/790).

Article 4 of the DSM Directive creates an exception regarding text and data mining (TDM). Works that are lawfully accessible can be reproduced and extracted for TDM purposes. However, if a copyright holder has expressly reserved the use in an appropriate manner, such as machine-readable means (opt-out), this exception does not apply. The robots.txt file is a representative means of such reservation.

robots.txt is a text file placed in the root directory of a website. It has been used since 1994 to instruct search engine crawlers not to

collect certain pages. The EU elevated this 30-year-old technical standard into a legal tool for copyright protection.

On August 1, 2025, the Code of Practice for GPAI approved by the EU concrete-ized this framework. The Code of Practice is structured around three pillars: transparency, copyright, and safety and security. GPAI model providers must respect reservation-of-rights indications such as robots.txt during the training data collection process, and retain model-related documentation for 10 years. Publishing a sufficiently detailed summary about the content used for training is an obligation directly imposed by Article 53, Paragraph 1 of the AI Act prior to the Code of Practice.

Compared to Korea, the difference is stark. The EU explicitly specified the TDM exception in copyright law while simultaneously guaranteeing the opt-out right. AI business operators can use content to which they have lawful access for training if the rights holder has not reserved rights via robots.txt or other means, and cannot use content that has been reserved. The rules are relatively clear. Korea has neither a TDM exception nor an opt-out system. It is a structure that leaves legal uncertainty for business operators and a gap in protection for copyright holders.

7. The Blank Cell

On May 7, 2025, Legislative Research Officer Baek Ji-yeon of the National Assembly Research Service published "Discussions and Legislative Tasks Regarding Disclosure of Training Data for Generative AI" (Issues and Perspectives No. 2355). This report points out that Korea lacks a legal basis requiring generative AI business operators to disclose the sources and composition of training data.

The EU made it mandatory for GPAI model providers to disclose a summary of training data. In the United States, pressure for disclosure practically operates through litigation. Japan has an explicit exception provision but is re-discussing its limitations. Korea defined "training data" in the AI Basic Act, but did not establish regulations regarding the collection methods or source disclosure of that data. Nor are there explicit provisions regarding reproduction for AI training in the Copyright Act.

The remaining legislative tasks are clear. Whether to establish explicit provisions in the Copyright Act regarding reproduction for AI training. Options include introducing a Japanese-style non-enjoyment purpose exception or bringing in an EU-style TDM exception and opt-out system. Whether to mandate AI business operators to disclose the sources and composition of training data.

Whether to guarantee copyright holders the right to check if their works were used for AI training and to refuse use. As of July 2026, none of these have been legislated.

A wide empty space lies between the definition of "training data" in Article 2, Subparagraph 12 of the AI Basic Act and the fair use provision of Article 35-5 of the Copyright Act. One side defined data broadly without mentioning copyright, while the other side set only fair use criteria without creating specific regulations regarding AI training.

When a judgment comes out in the broadcasters v. portal business operators lawsuit in the Seoul courts, it will effectively replace legislation. If fair use is broadly recognized, crawling will continue. If interpreted narrowly, licenses will be required for all training data. It is a situation where a single court ruling determines the data procurement cost for the entire industry.

When counsel for The New York Times wrote "billions of dollars" on the complaint, it was not merely about the amount of money. Shall AI training data be viewed as a free raw material, or as a subject of paid licensing? Korea's answer to that question has not yet come out.

Chapter 15. Intersection with the Personal Information Protection Act

One day in 2024, a data team at a major domestic platform company submitted an inquiry to its internal legal team: "We intend to use millions of posts publicly available on social media to train an AI model. Since these posts were set to public by users, do we need to obtain separate consent?" The legal team's answer consisted of two sentences: "Under the Personal Information Protection Act, consent is in principle required. However, while the AI Basic Act separately defines training data in Article 2, Subparagraph 12, it remains silent on the processing of personal information." Between these two sentences lies a practical gap in real-world compliance.

1. Relationship Between the Two Acts: General Act or Parallel Acts?

Article 5, Paragraph 1 of the AI Basic Act states: "Except as otherwise provided for specifically in other acts, artificial intelligence, etc. shall be governed by what is prescribed by this Act." Article 6, Paragraph 1 of the Personal Information Protection Act also shares nearly the same structure: "Except as otherwise provided for specifically in other acts, the processing and protection of personal information shall be governed by what is prescribed by this Act." Both acts claim themselves to be the general act.

Where does the issue of processing personal information contained in AI training data fall under this structure? The AI Basic Act defines training data as "data used for the development, utilization, etc. of artificial intelligence" (Article 2, Subparagraph 12). Although the statute itself recognizes that this data may contain personal information, it does not prescribe the requirements for the lawfulness of personal information processing.

The Personal Information Protection Act regulates the overall processing of personal information. Whether for AI training or statistical analysis, the act of collecting and using information that identifies a living individual falls under the scope of this Act. A report published by the Korea Legislation Research Institute in December 2025, titled "Measures to Enhance the Lawfulness of Personal Information Processing in Artificial Intelligence Development," analyzed that the AI Basic Act has effectively deferred the issue of personal information processing to the Personal Information Protection Act.

Article 34, Paragraph 3 of the AI Basic Act creates a single point of intersection: "Where an AI business operator has taken measures equivalent to each subparagraph of Paragraph 1 in accordance with

other statutes as prescribed by Executive Order [Enforcement Decree], the business operator shall be deemed to have taken the measures under Paragraph 1." A business operator that has fulfilled safety measures and obligations to guarantee the rights of data subjects under the Personal Information Protection Act is recognized as having fulfilled corresponding high-impact AI responsibilities within the scope of personal information processing and protection. While this serves as a mechanism to alleviate duplicate regulation, it is a one-way street. Compliance with the Personal Information Protection Act may satisfy obligations under the AI Basic Act, but compliance with the AI Basic Act does not exempt a business operator from obligations under the Personal Information Protection Act.

To summarize: with regard to the processing of personal information within AI training data, the Personal Information Protection Act serves as the general act, and the AI Basic Act provides no separate special provisions.

2. Personal Information in Training Data: The Issue of Consent for Collection and Use

Data is required to train AI models: text, images, voice, and video. If this data contains an individual's name, face, voice, location, or

behavioral patterns, it constitutes personal information. Article 2, Subparagraph 1 of the Personal Information Protection Act defines personal information as "information relating to a living individual that identifies the individual by name, resident registration number, image, etc." Even if an individual cannot be identified by the information alone, if it can be easily combined with other information to identify an individual, it is still personal information.

Article 15, Paragraph 1 of the Personal Information Protection Act enumerates seven lawful grounds for collection and use. These include the consent of the data subject (Subparagraph 1), special provisions in other acts (Subparagraph 2), performance of a contract (Subparagraph 4), and legitimate interest (Subparagraph 6). In large-scale data collection for AI training, obtaining individual consent from hundreds of thousands or millions of data subjects is practically impossible. Can a company that crawls text publicly available on the web to train a large language model obtain consent from every original author?

In practice, two grounds are discussed for collecting data without consent: legitimate interest under Article 15, Paragraph 1, Subparagraph 6, and use within a scope reasonably related to the

original purpose of collection under Article 15, Paragraph 3.

The legitimate interest provision has existed since the enactment of the law in 2011: "Where it is necessary to achieve the legitimate interest of the personal information controller, and where such interest manifestly takes priority over the rights of the data subject." The conditions are stringent. It must be substantially relevant to the legitimate interest and must not exceed a reasonable scope. There are not many situations where one can argue that the benefit of AI training "manifestly" takes priority over the privacy rights of data subjects.

Use within a reasonably related scope under Article 15, Paragraph 3 serves as the ground when using personal information that has already been lawfully collected for a purpose other than the original purpose. This may cover cases where data provided by users during service use is originally collected for service operation and later repurposed for training AI models. Factors such as whether disadvantage is caused to the data subject and whether measures to ensure safety have been taken must be considered.

Provision to a third party is even stricter. Article 17, Paragraph 1 of the Personal Information Protection Act in principle requires the

consent of the data subject for third-party provision. The act of an AI development business operator receiving training data from a company holding the data constitutes provision to a third party. As a ground for providing data without consent, there is provision within the scope of the purpose of collection (Article 17, Paragraph 1, Subparagraph 2), but whether "transferring data for AI training" falls within the scope of the original purpose of collection is almost always a subject of debate. Provision within a scope reasonably related to the original purpose of collection is separately grounded in Article 17, Paragraph 4. In this case as well, it must be considered whether disadvantage is caused to the data subject and whether measures to ensure safety, such as encryption, have been implemented.

3. The Pitfalls of Publicly Available Personal Information

Social media profiles, blog posts, YouTube videos, and employee introductions on corporate websites—this is information disclosed by data subjects themselves. Can this information be used for AI training?

The Personal Information Protection Act does not contain separate provisions for "publicly available personal information." As long as it is personal information, regardless of whether it is public, this Act applies. A data subject posting a text on social media is an act of

disclosing it for users of that platform to read, not an act of consenting to a third-party company using it as AI model training data.

In July 2024, the Personal Information Protection Commission issued the "Guidelines on Processing Publicly Available Personal Information for Artificial Intelligence (AI) Development and Services." It recognized that legitimate interest under Article 15, Paragraph 1, Subparagraph 6 could serve as a lawful basis when using publicly available personal information for AI training, while attaching conditions such as safety measures and the guarantee of data subjects' rights. An organizational structure was also established. The current Enforcement Rule on the Official Structure of the Personal Information Protection Commission includes an AI Privacy Team within the Personal Information Policy Bureau, a dedicated department handling personal information issues in the AI era.

The European Union took a clearer stance on this issue. Supervisory authorities in various countries have been issuing determinations on whether the legitimate interest provision under Article 6(1)(f) of the EU General Data Protection Regulation (GDPR) can be applied to AI training. In March 2023, the Italian Data Protection Authority

(Garante) temporarily suspended OpenAI's ChatGPT service in Italy, concluding that there was no lawful basis for processing. OpenAI subsequently cited legitimate interest as its basis, but the debate over whether web crawling for large language models can be justified by legitimate interest remains unresolved.

In Korean law, the conclusion points in the same direction. Even if it is publicly available personal information, it cannot be used directly for large-scale crawling for AI training purposes. Lawful grounds for collection, notification to data subjects, and safety measure obligations still follow.

4. Pseudonymized Information and Anonymized Information: Lawful Pathways for AI Training

Through the 2020 amendments to the Three Data Acts, the Personal Information Protection Act introduced the pseudonymized information system. This system opens up a lawful pathway for AI training data.

Pseudonymization refers to processing personal information by deleting or replacing part of it so that a specific individual cannot be identified without additional information (Article 2, Subparagraph 1-2). Pseudonymized information remains a subcategory of personal

information because the possibility of restoring it to its original state remains. However, Article 28-2, Paragraph 1 grants a special exception to pseudonymized information: "A personal information controller may process pseudonymized information without the consent of data subjects for purposes such as compilation of statistics, scientific research, and preservation of records in the public interest."

Does AI model training qualify as "scientific research"? Article 2, Subparagraph 8 of the Personal Information Protection Act defines scientific research as "research applying scientific methods, such as technology development and demonstration, basic research, applied research, and privately funded research." There is ample room to view AI model development by private companies as falling under technology development and demonstration. In the National Assembly discussions during the 2020 amendment, a broad definition encompassing industrial research was adopted to revitalize the data economy.

While pseudonymized information can be processed without consent, there are conditions. It must not be processed for the purpose of identifying a specific individual (Article 28-5, Paragraph 1).

If information that can identify an individual is generated during processing, it must be suspended and destroyed immediately (Paragraph 2 of the same Article). Safety measure obligations are also imposed on pseudonymized information (Article 28-4). Combining pseudonymized information between different personal information controllers must be conducted through a specialized institution designated by the Personal Information Protection Commission or the relevant central administrative agency (Article 28-3).

The limitations of the pseudonymized information system are also clear. Large language models train on entire texts. Pseudonymizing only names does not resolve the issue. If descriptions that can identify an individual remain within the context, that text is still personal information.

Anonymized information is different. Article 58-2 of the Personal Information Protection Act states that this Act shall not apply to "information that can no longer identify an individual even if combined with other information, reasonably taking into account time, cost, technology, etc." Completely anonymized data is not personal information. Because it is not subject to the Personal Information Protection Act, there is no requirement for consent,

safety measures, or processing restrictions. It can be used for AI training without restriction. The problem is that the standard for "complete anonymization" is high. Technologies capable of re-identifying individuals within large-scale datasets continue to advance.

A realistic pathway for a business operator is summarized as follows: to use training data containing personal information without consent, one must either process it for scientific research purposes after pseudonymization or completely anonymize it to fall outside the scope of application of the Personal Information Protection Act altogether. Either way, technical measures and legal verification are required side by side.

5. Two Rights to Demand Explanation

Both the AI Basic Act and the Personal Information Protection Act prescribe a "right to receive an explanation." However, their statutory structures and scopes differ.

Article 3, Paragraph 2 of the AI Basic Act reads: "An affected person shall be entitled to receive clear and meaningful explanations regarding the primary criteria, principles, etc. utilized to derive the final result of artificial intelligence, to the extent technically and

reasonably feasible." This provision is placed under basic principles. It does not prescribe specific procedures or sanctions for non-compliance. An "affected person" refers to a person whose life, physical safety, or fundamental rights are significantly affected by an artificial intelligence product or service (Article 2, Subparagraph 9). It applies not to all AI users, but only to persons who are significantly affected.

Article 37-2 of the Personal Information Protection Act is different. This provision was newly introduced through the 2023 amendment and came into effect on March 15, 2024: "Where a decision made by processing personal information through a fully automated system (including systems applying artificial intelligence technology), excluding automated dispositions by administrative agencies under Article 20 of the Administrative Procedures Act, significantly affects their rights or obligations, a data subject shall have the right to refuse such decision toward the relevant personal information controller." It grants both the right to demand explanation (Paragraph 2) and the right to refuse (Paragraph 1).

Comparing the differences makes the outline clear.

The right to demand explanation under the AI Basic Act is granted to an "affected person." It targets AI systems in general. Regardless of whether it is high-impact or not, any person significantly affected qualifies. However, remaining as a basic principle provision, it lacks implementation procedures. On the other hand, Article 37-2 of the Personal Information Protection Act is granted to a "data subject." It targets "decisions made by processing personal information through a fully automated system." A data subject can refuse the decision and demand an explanation, and the personal information controller must take necessary measures, such as re-processing with human intervention, unless there is a justifiable reason (Paragraph 3). There is also an obligation to disclose the criteria and procedures for automated decisions as well as how personal information is processed (Paragraph 4).

Suppose a person is denied a loan by an AI loan screening system. This individual is both an affected person under the AI Basic Act and a data subject under the Personal Information Protection Act. The rights to demand explanation under both acts apply simultaneously. In practice, it remains unclear under which act and at what level a business operator must provide an explanation, or which side takes precedence if requirements under the two acts conflict. Because

Article 37-2 of the Personal Information Protection Act is backed by implementation procedures and sanctions, its practical binding force is stronger.

6. Profiling and Automated Decision-Making

Profiling refers to the act of automatically processing personal information to analyze or predict specific aspects of an individual. Credit score calculation, customer behavior analysis, and targeted advertising fall under this category. This is precisely what AI does.

"Decisions made by processing personal information through a fully automated system" under Article 37-2 of the Personal Information Protection Act encompasses the outcomes of AI-based profiling. Provisos are attached to this provision: the right to refuse does not apply where the automated decision has obtained the consent of the data subject (Article 15, Paragraph 1, Subparagraph 1), where it is inevitable to comply with statutory obligations (Subparagraph 2), or where it is necessary for the performance of a contract (Subparagraph 4).

This structure resembles Article 22 of the EU GDPR. Article 22(1) of the GDPR stipulates that "The data subject shall have the right not to be subject to a decision based solely on automated processing,

including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her." It is a prohibition in principle. The only exceptions are when necessary for entering into, or performance of, a contract, authorized by Union or Member State law, or based on the data subject's explicit consent (Article 22(2)).

The difference between Korean law and the GDPR lies in their starting points. The GDPR prohibits automated decisions in principle and opens up exceptions. Korea's Personal Information Protection Act does not prohibit automated decisions, but instead grants data subjects the right to refuse and the right to demand explanation. The approaches differ: the EU follows "prohibition in principle + exceptional permission," whereas Korea follows "permission + right to refuse + right to explanation."

The AI Basic Act does not use the term profiling. However, Article 2, Subparagraph 4, Item (g) enumerates "judgments or evaluations that significantly affect an individual's rights and obligations, such as hiring and loan screening," as application areas of high-impact AI. Recruitment screening, loan screening, and insurance risk assessment—when AI analyzes an individual's data to make judgments in these areas, the high-impact AI obligations under the AI

Basic Act (risk management, explanation, human intervention, and documentation) and the automated decision-making regulations under the Personal Information Protection Act apply simultaneously.

Suppose a financial company conducts loan screening using AI. Under the AI Basic Act, it must establish risk management measures, explain key criteria, and maintain a system managed and supervised by humans. Under the Personal Information Protection Act, it must disclose the criteria and procedures for automated decisions, and if a rejected applicant demands an explanation, it must re-process the application with human intervention. Provisions regarding automated evaluation under the Use and Protection of Credit Information Act overlap as well. Three statutes overlap on a single service.

7. Two Impact Assessments

Article 35 of the AI Basic Act prescribes high-impact AI impact assessment: "Where an AI business operator provides a product or service using high-impact artificial intelligence, it shall endeavor in advance to assess the impact on human fundamental rights. In this case, the impact assessment shall ensure that the characteristics of vulnerable groups to artificial intelligence are reflected, taking into account the nature of the product or service using high-impact

artificial intelligence." This is a voluntary duty of effort. For state agencies to use high-impact AI products, they must give priority consideration to products for which an impact assessment has been conducted (Paragraph 2).

Article 33 of the Personal Information Protection Act prescribes personal information impact assessment. Where the operation of personal information files meeting the criteria prescribed by Executive Order [Enforcement Decree] is likely to infringe upon the personal information of data subjects, the head of a public institution must conduct an impact assessment and submit the results to the Protection Commission. This is mandatory for public institutions. Private personal information controllers bear a voluntary duty of effort to "actively endeavor" (Paragraph 11).

While the two impact assessments have similar names, their subjects of assessment differ. The AI impact assessment examines the impact of the AI system on human fundamental rights. AI-specific risks such as bias, discrimination, safety, and transparency are evaluation items. The personal information impact assessment examines the impact of operating personal information files on the infringement of data subjects' personal information. Items evaluated

include the volume of personal information processed, whether it is provided to third parties, and the likelihood of infringing data subjects' rights.

If a high-impact AI processes personal information, a public institution may need to conduct both impact assessments. A public institution adopting an AI recruitment system must perform an AI impact assessment under the AI Basic Act and a personal information impact assessment under the Personal Information Protection Act, respectively. There are overlapping evaluation items (discrimination and bias, infringement of data subjects' rights) as well as non-overlapping items (technical safety of AI, retention period of personal information).

The duplicate regulation relief provision under Article 34, Paragraph 3 of the AI Basic Act does not extend to this. This deemed compliance applies only to business operator responsibilities under Article 34, Paragraph 1, while the impact assessment under Article 35 is a separate statutory provision. Even if a personal information impact assessment has been completed, AI-inherent risks (bias, opacity, vulnerability to adversarial attacks) are not covered by that evaluation, leaving the AI impact assessment as a separate task of

effort.

8. Distance from Article 22 of the EU GDPR

Article 22 of the EU GDPR is a preceding norm that influenced Korean law. When Article 37-2 of the Personal Information Protection Act was introduced through the 2023 amendment, Article 22 of the GDPR was examined as a reference model in legislative discussions. However, there is a gap between the structures and effects of the two provisions.

Article 22 of the GDPR targets "decisions based solely on automated processing." Decisions involving meaningful human intervention are excluded from its scope. Article 37-2 of Korea's Personal Information Protection Act targets "decisions made by processing personal information through a fully automated system." The term "fully" corresponds to "solely" in the GDPR, but room for interpretation remains on both sides regarding what degree of human involvement removes a process from "full automation."

The EU AI Act built an additional layer on top of Article 22 of the GDPR. High-risk AI systems in Annex III cover AI used in domains such as employment, education, essential services, and access to credit, imposing mandatory conformity assessments, technical

documentation, log keeping, and human oversight for such AI. In areas where the right to refuse under Article 22 of the GDPR overlaps with conformity assessments under the AI Act, EU law applies both statutes concurrently. Korea shares the same structure: Article 37-2 of the Personal Information Protection Act and Article 34 of the AI Basic Act apply concurrently.

The point where the difference becomes stark is the level of sanctions. The penalty fine for GDPR violations is up to EUR 20 million or 4% of total worldwide annual turnover. While Korea's Personal Information Protection Act contains administrative fine provisions (Article 64-2), the maximum administrative fine under the AI Basic Act is 30 million KRW. This gap directly affects a business operator's motivation to comply.

9. Practical Checklist: Five Things Business Operators Must Verify

Here is a summary of items that business operators developing or using AI must verify at the intersection of the Personal Information Protection Act and the AI Basic Act.

Does the training data contain personal information? If text, image, voice, or video data contains information identifying an individual, the Personal Information Protection Act applies. Unconsented

collection without pseudonymization or anonymization carries a high risk of unlawfulness.

Has pseudonymization or anonymization been performed?

Pseudonymized information can be processed without consent for scientific research purposes, but is accompanied by prohibitions on re-identification and safety measure obligations. Anonymized information is not subject to the Personal Information Protection Act, but the burden of proving the sufficiency of anonymization lies with the business operator.

Does the AI system make automated decisions? If a system makes automated decisions that significantly affect an individual's rights or obligations, the right to refuse and the right to demand explanation under Article 37-2 of the Personal Information Protection Act, as well as the right to demand explanation under Article 3, Paragraph 2 of the AI Basic Act, apply simultaneously. A system for human intervention must be established.

Does high-impact AI process personal information? If personal information is processed in high-impact AI domains (recruitment, loan screening, healthcare, public services, etc.), high-impact AI obligations under the AI Basic Act overlap with safety measures and

impact assessment obligations under the Personal Information Protection Act. While compliance with the Personal Information Protection Act may satisfy certain obligations under the AI Basic Act pursuant to Article 34, Paragraph 3 of the AI Basic Act, AI-specific obligations (risk management, transparency) must be fulfilled separately.

Is publicly available personal information crawled and used for training? Even if information was disclosed by data subjects, large-scale collection for AI training purposes requires separate lawful grounds. Even when invoking the legitimate interest provision (Article 15, Paragraph 1, Subparagraph 6), proving that such interest manifestly takes priority over the rights of data subjects is not easy.

10. Blank Spaces in the Law

The AI Basic Act remained silent on personal information processing, and the Personal Information Protection Act has yet to establish special provisions tailored to the new reality of AI training. While the pseudonymized information system offers one pathway, in the face of the reality where large language models train on the entire web, the path of "use for scientific research purposes after pseudonymization" is narrow.

The Personal Information Protection Commission's issuance of guidelines and operation of the AI Privacy Team represent efforts to fill this blank space. The amended Personal Information Protection Act promulgated on March 10, 2026 (Act No. 21445, scheduled to take effect on September 11, 2026, and July 1, 2027) also does not appear to contain special provisions directly addressing AI training. Statutory special provisions tailored to the AI era remain a task for future amendments.

Article 5, Paragraph 2 of the AI Basic Act states: "Where other acts are enacted or amended regarding artificial intelligence, etc., they shall conform to the purpose of this Act." Subsequent AI-related amendments to the Personal Information Protection Act will also be bound by this provision. How the two acts will synthesize will be decided not by statutory text, but by legal practice and enforcement. For now, this is a time when questions from business operators continue to pile up between the two acts.

Chapter 16. Overlap with Industry-Specific Adjacent Regulations

In March 2026, a compliance officer at a startup operating an AI-based skin analysis app pulled up an Excel file during a legal team meeting. The vertical axis listed twelve features of the company's service, and the horizontal axis listed six laws: the AI Basic Act, the Medical Devices Act, the Digital Medical Products Act, the Personal Information Protection Act, the Act on Fair Labeling and Advertising, and the Electronic Commerce Act. The cells were densely packed with numbers of mandatory clauses. Risk management obligations alone overlapped across three laws, and user notification obligations were required separately by four laws. The officer asked, "Which of these laws must we comply with first?"

The AI Basic Act is a horizontal general law. It imposes obligations on all business operators using AI regardless of the industry sector. However, South Korea already has vertical regulations accumulated over decades in sectors such as finance, healthcare, telecommunications, personal information, and consumer protection. When a new horizontal law overlaps on top of existing vertical regulations, business operators are placed in a position where they must prepare similar documents in different formats for

two or more entities regarding the same act. The National Assembly Research Service (NARS) pointed out even before the enforcement of the AI Basic Act that sector-specific AI legislative discussions were necessary in areas such as finance, healthcare, and robotics, and the Korea Legislation Research Institute (KLRI) also warned that the hierarchy was reversed because the Basic Act was enacted first before sector-specific AI regulations could mature.

1. The Principle of Article 5: Special Law Prevails

Article 5 of the AI Basic Act consists of two paragraphs.

Paragraph 1 declares, "Except as otherwise specially provided in other laws regarding artificial intelligence, etc., this Act shall apply." It is a classic provision giving priority to special laws over general laws. If the Medical Devices Act contains separate provisions for AI medical devices, the Medical Devices Act applies first in that domain, and the AI Basic Act steps back.

Paragraph 2 stipulates, "When enacting or amending other laws regarding artificial intelligence, etc., such enactments or amendments shall conform to the purposes of this Act." This means that even if sector-specific laws are enacted later, they should not deviate from the philosophy and direction of the AI Basic Act.

The principle is clear. If existing sector-specific laws already contain provisions governing AI, those laws take priority. If not, the AI Basic Act applies supplementally. However, there is a gap between principle and practice. This is because interpretations can diverge on what specifically constitutes "cases where special provisions exist," or whether general safety obligations under existing laws qualify as "special provisions" capable of substituting for high-impact AI obligations under the AI Basic Act.

2. Enforcement Decree Annex 1: A Tool to Cut Overlapping Burdens

The AI Basic Act has established a concrete solution for the issue of overlap: Article 34, Paragraph 3. "Where an artificial intelligence business operator has implemented measures equivalent to any of the items under Paragraph 1 pursuant to other statutes as prescribed by Presidential Decree, the business operator shall be deemed to have implemented the measures under Paragraph 1." This means that complying with another law will be recognized as compliance with this Act as well.

Article 27, Paragraph 5 of the Enforcement Decree is that Presidential Decree. "Cases in which an artificial intelligence business operator is deemed to have implemented the measures under Article 34, Paragraph 1 of the Act pursuant to Paragraph 3 of the same Article

shall be as specified in Annex 1." Annex 1 lists item by item which of the five obligations under Article 34, Paragraph 1, Items 1 through 5 of the AI Basic Act (risk management, explanation, user protection, human supervision, and documentation) are deemed to be fulfilled when a business operator complies with specific provisions of specific laws in each sector. Although Article 34, Paragraph 1 also contains Item 6, which requires the implementation of matters deliberated and resolved by the Committee, Annex 1 targets Items 1 through 5 for recognition of compliance.

Eight categories are listed in Annex 1: digital medical devices, autonomous vehicles, cybersecurity of motor vehicles under the Motor Vehicle Management Act, maritime autonomous surface ships, nuclear materials and nuclear facilities, loan screening, safety protection measures under the Framework Act on Intelligent Information Technology, and protection measures under the Personal Information Protection Act.

Digital medical devices dominate in terms of the number of items. If a business operator establishes a quality management system pursuant to the Digital Medical Products Act and obtains a determination of conformity with quality management standards, it

is deemed to have implemented the risk management obligation (Item 1) under the AI Basic Act. If manufacturing authorization or import authorization is obtained and information is displayed on the software, the explanation obligation (Item 2) is also satisfied.

Complying with compliance matters fulfills user protection (Item 3), designating a quality manager satisfies human supervision (Item 4), and establishing a quality management system along with obtaining conformity determinations fulfills documentation (Item 5)—all five obligations can be resolved within the framework of the Digital Medical Products Act.

The field of loan screening is also included. Complying with the explanation obligation pursuant to Article 35-2 of the Credit Information Act and establishing automated evaluation explanation procedures pursuant to Article 36-2, Paragraph 1 is deemed compliance with the explanation obligation under the AI Basic Act. Fulfilling duties under Article 10 of the Financial Consumer Protection Act also satisfies the user protection obligation.

Autonomous vehicles are deemed to have fulfilled the obligations of risk management (Item 1), user protection (Item 3), and documentation (Item 5) if performance certification and modification

certification are obtained under the Act on Promotion and Support for Commercialization of Autonomous Vehicles. Human supervision (Item 4) is recognized only when all duties under Article 43 of the same Act are fulfilled, and there is no recognized item for the explanation obligation (Item 2). Maritime autonomous surface ships are deemed to have fulfilled all five obligations if safety assessments and operation approvals are obtained under the Maritime Autonomous Surface Ships Act.

Motor vehicles under the Motor Vehicle Management Act are also included in Annex 1. If certification and modification certification for the automotive cybersecurity management system are obtained and all duties under Article 34-5 of the same Act are fulfilled, the risk management obligation (Item 1) is deemed to be implemented. For nuclear materials and nuclear facilities, if relevant details are included in compliance matters or the code of ethics under the Act on the Management and Supervision of Nuclear Power Business Operators, etc. to Prevent Malpractice in Nuclear Power Industry, the obligations of explanation (Item 2) and human supervision (Item 4) are deemed to be implemented.

The Personal Information Protection Act is handled as a separate item (Item 7 of Annex 1). If measures and obligations under 12 provisions—such as Article 29 (Duty to Take Safety Measures), Article 30 (Disclosure of Personal Information Processing Policy), Article 31 (Designation of Personal Information Protection Officer), Article 35 (Inspection of Personal Information), Article 36 (Correction and Erasure of Personal Information), and Article 37 (Suspension of Processing of Personal Information) of the Personal Information Protection Act—are implemented, the business operator is deemed to have implemented measures under each item of Article 34, Paragraph 1 of the AI Basic Act "with respect to the processing and protection of personal information." Because the qualifying phrase "with respect to the processing and protection of personal information" is attached, risk management obligations unrelated to personal information are not exempted.

Implementing safety protection measures under Article 60, Paragraph 1, Items 1 through 3 and Item 5 of the Framework Act on Intelligent Information Technology recognizes compliance with the risk management (Item 1) and explanation (Item 2) obligations, while implementing protection measures under Item 4 of the same paragraph recognizes compliance with the human supervision (Item

4) obligation.

3. The Places Left Empty by Annex 1

There are domains not named in Annex 1.

Energy supply, drinking water production, healthcare system establishment, nuclear safety (only partially included in Annex 1), biometric information analysis, public institution decision-making, and student evaluation. Among the 11 high-impact AI domains listed in Article 2, Item 4 of the AI Basic Act, the areas where Annex 1 has established concrete criteria for compliance recognition are limited to medical devices, loan screening, transportation (autonomous vehicles, motor vehicle cybersecurity, maritime autonomous surface ships), and nuclear power (partially). Business operators conducting business in the remaining domains must comply with both existing statutes and the AI Basic Act.

The recruitment field is a prime example. The Fair Recruitment Procedure Act does not contain separate provisions regarding AI interviews or AI document screening. Although the Ministry of Employment and Labor released the 'AI Recruitment System Guidelines' in December 2023, they lack legal binding force. Business operators using AI for recruitment must fully implement the

high-impact AI obligations of the AI Basic Act while simultaneously complying with existing obligations under the Fair Recruitment Procedure Act. Because recruitment is omitted from Annex 1, complying with the Fair Recruitment Procedure Act does not exempt a business operator from obligations under the AI Basic Act.

Student evaluation is similar. The Framework Act on Education contains no provisions regarding AI-based evaluation. The high-impact AI obligations of the AI Basic Act apply directly to AI-assisted grading of the College Scholastic Ability Test or calculation of student grades, but until the Ministry of Education issues separate guidelines or public notices, the methods of implementation remain ambiguous.

4. Act on Fair Labeling and Advertising: Two Laws Applied to a Single Advertisement

Consider a situation where an advertising image generated by AI is posted on the Internet.

Article 31, Paragraph 2 of the AI Basic Act requires that output created by generative artificial intelligence display "the fact that it was generated by generative artificial intelligence." Paragraph 3 instructs that virtual images or videos that are difficult to distinguish

from reality be notified or displayed in a "manner clearly discernible by users."

Article 3, Paragraph 1 of the Act on Fair Labeling and Advertising (Fair Labeling and Advertising Act) prohibits labeling and advertising that "is likely to deceive or mislead consumers." False, exaggerated, or deceptive labeling and advertising fall into this category. If a cosmetics advertisement uses an AI-generated model image while portraying it as actual usage results, it may constitute deceptive advertising under the Fair Labeling and Advertising Act.

The obligations under the two laws overlap. The AI Basic Act requires displaying the fact itself that "it was generated by AI," whereas the Fair Labeling and Advertising Act requires "do not deceive consumers." Although the direction is aligned, the methods of display differ. Fulfilling the display obligation under the AI Basic Act does not automatically satisfy the obligation prohibiting unfair advertising under the Fair Labeling and Advertising Act. The reverse is also true. Even if a notice states "This image was generated by AI," if the content of the advertisement itself is an exaggeration, it constitutes a violation of the Fair Labeling and Advertising Act.

Applying the logic of Article 5, since the Fair Labeling and Advertising Act constitutes "special provisions" regarding the truthfulness of advertising, the Fair Labeling and Advertising Act takes priority in judging the deceptive nature of advertising content. The obligation to display AI-generated output under the AI Basic Act operates on a layer distinct from the Fair Labeling and Advertising Act. Ultimately, business operators must comply with both laws. The Fair Labeling and Advertising Act is not included in Annex 1.

In a situation where the Fair Trade Commission has not yet issued examination guidelines or advice regarding AI-generated advertising, concrete criteria concerning the method, location, and size of displays are not sufficiently defined in either law. Business operators have no choice but to lay the transparency guidelines of the AI Basic Act side by side with the existing examination guidelines of the Fair Labeling and Advertising Act and reconcile the overlapping parts themselves.

5. Digital Medical Products Act: A Tightly Connected Domain

The law that gears together precisely like cogs with the AI Basic Act is the Digital Medical Products Act.

The Digital Medical Products Act was promulgated on January 23, 2024, and took effect on January 24, 2025, with certain provisions taking effect on January 24, 2026. Article 2, Item 4, Subitem (d) of the AI Basic Act specifically names the Digital Medical Products Act. It explicitly specifies that "development and use of medical devices under the Medical Devices Act and digital medical devices under the Digital Medical Products Act" is one domain of high-impact AI. The portion of Subitem (d) concerning digital medical devices came into effect on January 24, 2026, entering into force with a two-day gap following the enforcement date of the AI Basic Act on January 22, 2026.

Let us look at the obligations of a business operator creating AI-based medical devices. The Digital Medical Products Act, under the jurisdiction of the Ministry of Food and Drug Safety (MFDS), regulates the entire process from manufacturing authorization, quality management systems, clinical trials, software information display, designation of quality managers, to post-market safety management. The AI Basic Act requires risk management, explanation, user protection, human supervision, and documentation.

Thanks to Annex 1, a business operator that faithfully adheres to the Digital Medical Products Act framework does not need to additionally perform high-impact AI obligations under the AI Basic Act. This is because corresponding provisions under the Digital Medical Products Act are designated for all five obligation items in Annex 1.

However, there is one point of caution. What Annex 1 exempts is limited to the obligations under Article 34 (Duties of High-Impact AI Business Operators) of the AI Basic Act. Prior notification and output labeling under Article 31 (Duty to Ensure Transparency) and best-effort obligations under Article 35 (Impact Assessment) are not included within the scope of Annex 1. Even an AI medical device business operator must separately provide prior notification to users stating, "This service is operated based on artificial intelligence."

The EU AI Act takes a different approach to medical devices. The EU Medical Devices Regulation (MDR) and In Vitro Diagnostic Medical Devices Regulation (IVDR) are product safety regulations listed in Annex I of the AI Act. Medical devices that undergo conformity assessment pursuant to these regulations are automatically classified as high-risk AI under Article 6(1) of the AI Act, and the conformity assessment procedure is integrated into one. South Korea adopts a

mutual recognition approach rather than integration. While compliance with existing statutes is recognized as fulfillment, the two laws remain separate frameworks.

6. Finance: Between Guidelines and Laws

The AI regulatory landscape in the financial sector is complex.

The Financial Services Commission (FSC) announced the 'AI Guidelines for the Financial Sector' in July 2021 and introduced the 7 Principles for Financial AI in December 2024. These guidelines are administrative guidance without legal binding force. They recommend fairness, explainability, and data quality management for AI credit scoring models.

Article 36-2 of the Credit Information Act regulates by statute the obligation to explain automated evaluations and the right to object. An individual credit information subject can receive explanations regarding the results, major criteria, and overview of basic information of automated evaluations, and if the basic information is inaccurate, may request correction or deletion, or request recalculation of results. This provision existed even before the AI Basic Act was enacted.

Article 34, Paragraph 1, Item 2 of the AI Basic Act directs high-impact AI business operators to establish and implement measures to explain "major criteria utilized in deriving final results," among other things. A financial company using AI for loan screening becomes subject to both the explanation obligation under the Credit Information Act and the explanation obligation under the AI Basic Act simultaneously.

Annex 1 resolves this issue. Complying with the explanation obligation under Article 35-2 of the Credit Information Act and maintaining procedures under Article 36-2, Paragraph 1 is deemed implementation of the explanation obligation under the AI Basic Act. Fulfilling duties under Article 10 of the Financial Consumer Protection Act also satisfies the user protection obligation.

The problem lies in financial AI outside loan screening: insurance underwriting screening, fraud detection systems (FDS), asset management robo-advisors, and internal credit rating calculations. While these may fall under Article 2, Item 4, Subitem (g) of the AI Basic Act—"judgment or evaluation significantly affecting the rights and duties of individuals, such as recruitment and loan screening"—Annex 1 specifically points only to "loan screening." It is

unclear whether financial companies using AI in insurance or investment are automatically exempted from AI Basic Act obligations simply by complying with existing financial regulations.

Merely following the FSC's AI guidelines is not recognized as fulfillment of AI Basic Act obligations. This is because Annex 1 takes only obligations at the statute and Enforcement Decree level as targets of mutual recognition, while guidelines are administrative guidance without legal binding force. The dual burden on financial companies to simultaneously comply with the FSC guidelines and the obligations of the AI Basic Act remains intact.

7. Personal Information Protection Act and Credit Information Act: Where Overlapping Layers Are Thickest

Almost without exception, AI systems handle personal information. Personal information is included in training data, and inference results contain judgments regarding individuals. The AI Basic Act and the Personal Information Protection Act apply simultaneously to virtually all AI business operators.

Item 7 of Annex 1 attempted to handle this issue comprehensively. It enumerates 12 provisions of the Personal Information Protection Act and provides that if these are implemented, the business operator

is deemed to have implemented the measures under each item of Article 34, Paragraph 1 of the AI Basic Act "with respect to the processing and protection of personal information." Included are duties to take safety measures (Article 29), disclosure of processing policies (Article 30), designation of protection officers (Article 31), right to inspection (Article 35), right to correction and erasure (Article 36), and right to suspension of processing (Article 37).

The qualifying phrase "with respect to the processing and protection of personal information" is the key. Among the risks of AI systems, only the portion related to personal information is exempted. For risks not directly tied to personal information—such as discrimination due to algorithmic bias, infringement of fundamental rights caused by unexplainable judgments, or safety accident risks—the obligations under the AI Basic Act remain fully intact.

Professor Do-seung Kim's report published by the Korea Legislation Research Institute in December 2025, titled "Measures to Enhance the Legality of Personal Information Processing in Artificial Intelligence Development," directly analyzes this point of conflict. The core observation is that a gap exists between the consent principle under the Personal Information Protection Act and the reality of AI

development during the training data collection stage. Training AI on tens of thousands of medical images often creates situations that are practically impossible under the consent framework of the Personal Information Protection Act. Although the AI Basic Act contains separate provisions on training data, its relationship with the Personal Information Protection Act has not yet been fully resolved.

The Credit Information Act serves as a special law to the Personal Information Protection Act in the financial sector. It contains automated evaluation explanation obligations (Article 36-2) as well as separate provisions for pseudonymous and anonymous processing. A financial company conducting AI loan screening stands at the intersection of three laws: the Credit Information Act, the Personal Information Protection Act, and the AI Basic Act. While Annex 1 allows the explanation obligation under the Credit Information Act to substitute for the explanation obligation under the AI Basic Act, obligations under the Personal Information Protection Act regarding consent, inspection, correction, and erasure still apply on a separate layer.

8. How Far Do Anti-Duplication Burden Mechanisms Function?

Let us reexamine the structure of Annex 1.

What Annex 1 exempts is confined to the five obligations under Article 34, Paragraph 1 of the AI Basic Act (risk management, explanation, user protection, human supervision, and documentation). Prior notification and output display under Article 31 (Duty to Ensure Transparency), high-performance AI risk management and result submission under Article 32 (Duty to Ensure Safety), preliminary review under Article 33 (Confirmation of High-Impact AI), and best-effort obligations under Article 35 (Impact Assessment) are not subject to Annex 1. In other words, even an AI medical device business operator that perfectly complies with the Digital Medical Products Act must independently fulfill transparency notification obligations and impact assessment best-effort obligations under the AI Basic Act.

Article 27, Paragraph 3 of the Enforcement Decree is another mechanism reducing overlapping burdens. Where an upstream artificial intelligence business operator has implemented measures under Article 34, Paragraph 1, Items 1 through 3 (risk management, explanation, user protection), a deploying business operator that receives the AI system is deemed to have implemented the same measures "unless it has made significant functional modifications such as significantly altering the original purpose or use." If the

upstream artificial intelligence business operator has already conducted risk management, the deploying business operator does not need to repeat it. However, if the deploying business operator applies the AI system to a different context and introduces meaningful modifications, this exemption does not apply.

9. NARS's Warning: Sector-Specific Legislation Is Necessary

In "Legislative Tasks for Supplementing the AI Basic Act Prior to Enforcement," published in March 2025, Legislative Research Officer Jun-hwa Jung of the National Assembly Research Service presented tasks that needed addressing before enforcement. One of them was that sector-specific AI legislative discussions in fields such as finance, healthcare, and robotics were necessary.

Research Fellow Won-jun Jung's September 2025 report from the Korea Legislation Research Institute, titled "A Study on Rational Regulation and Development Plans for the AI Basic Act," points out a more fundamental problem. Although the AI Basic Act bears the name "Basic Act," the original function of a basic act is to set the direction for subordinate individual laws. The customary sequence is that individual laws are enacted first and a basic act later encompasses them; however, South Korea enacted the Basic Act first when sector-specific AI laws barely existed. Because the Basic Act

directly imposes obligations on business operators while individual laws remain absent, points of awkward overlap with existing sector-specific regulations arise.

Research Fellow Won-jun Jung recommends that competent ministries for each sub-domain of high-impact AI overhaul AI-related regulations in their respective fields and explicitly clarify their relationship with the AI Basic Act. He argues that the current Annex 1 is merely a start and should be expanded to domains not yet included in the annex, such as recruitment, public services, education, energy, and biometrics.

10. The EU's Approach: Embedding AI into Existing Regulations

The EU AI Act reflected its relationship with existing regulations into its design from the outset.

When AI is embedded as a safety component in products subject to the EU's product safety harmonization legislation listed in Annex I—such as medical devices (MDR/IVDR), machinery (Machinery Regulation), motor vehicle general safety (General Safety Regulation), toy safety, and marine equipment—the high-risk AI requirements of the AI Act are evaluated together within the existing conformity assessment procedure for that product. Rather than evaluating twice,

it is integrated into a single procedure.

South Korea did not adopt this integrated model. Authorization procedures under the Digital Medical Products Act and implementation of obligations under the AI Basic Act exist as separate procedures, but fulfilling the former is "recognized" as fulfilling the latter. While the practical outcome may be similar, the legal structure is different. The EU finishes with a single certificate, whereas South Korea requires verifying compliance with each of the two laws separately.

The EU's Digital Omnibus amendment pushed this integration even further. By postponing the application deadlines for high-risk AI obligations to December 2027 (for standalone AI) and August 2028 (for embedded AI), it bought time to secure alignment with existing product safety regulations. It also set a target to reduce regulatory costs for small and medium-sized enterprises by 35%.

China is different again. Because China enacted sector-specific regulations first without a comprehensive AI law, there is inherently less room for overlapping issues to arise. This is because the Provisions on the Management of Algorithmic Recommendations, the Provisions on the Management of Deep Synthesis, and the Interim

Measures for the Management of Generative AI Services narrowly specify their respective scopes of application. However, there is a cost to this approach: there are no general principles establishing which regulation takes priority when a single AI service falls under multiple regulations simultaneously. One of the reasons China is preparing a comprehensive artificial intelligence law is precisely to integrate this fragmented regulatory framework.

11. Order of Evaluation Business Operators Should Follow in Practice

To business operators asking which law to comply with first, we suggest the following order.

First, check which sector-specific regulation your company's AI service already falls under: whether it is a product with medical device authorization, a service operated under financial business licensing, or a vehicle with autonomous driving performance certification. If sector-specific regulation exists, that regulation is the primary target for compliance. Pursuant to Article 5, Paragraph 1 of the AI Basic Act, "where special provisions exist in other laws," those laws take priority.

Next, check Annex 1 of the Enforcement Decree. If the provisions of sector-specific laws implemented by your company are listed in

Annex 1, you do not need to additionally perform the corresponding obligations under Article 34 of the AI Basic Act. If a quality management system under the Digital Medical Products Act is established, the risk management obligation under the AI Basic Act is satisfied.

Next, check obligations under the AI Basic Act that are not included in Annex 1. Prior notification for transparency (Article 31), confirmation of high-impact AI (Article 33), and impact assessment best-effort obligations (Article 35) must be implemented regardless of Annex 1.

Finally, if your company's AI service falls under the 11 high-impact AI domains but belongs to a sector not included in Annex 1 (such as recruitment, public services, student evaluation, energy, etc.), you must fully perform the obligations under Article 34 of the AI Basic Act. Complying with existing sector-specific laws does not grant an exemption.

Even following this sequence, ambiguous cases will arise. In such cases, inquiries can be made to the AI Basic Act Support Desk. The fact that 552 cases were received by the 10th week of enforcement, with the top items among online inquiries being "transparency"

(56.4%) and "confirmation of high-impact AI" (17%), demonstrates where the confusion experienced by business operators is concentrated.

12. For the Regulatory Map to Be Completed

Annex 1 of the Enforcement Decree of the AI Basic Act is the first edition of preventing overlapping burdens. Digital medical devices, autonomous vehicles, motor vehicle cybersecurity, maritime autonomous surface ships, nuclear materials and nuclear facilities, loan screening, the Framework Act on Intelligent Information Technology, and the Personal Information Protection Act—eight categories are listed. The rest remain blank.

The expansion of Annex 1 was not included even in the amended Enforcement Decree promulgated on July 20 following legislative notice on May 21, 2026. The work remains for competent sector-specific ministries—such as the Financial Services Commission, the Ministry of Employment and Labor, the Ministry of Education, the Ministry of Environment, and the Ministry of Land, Infrastructure and Transport—to overhaul AI regulations in their respective domains and officially clarify their relationship with the AI Basic Act. Until that work is completed, business operators must make decisions themselves at the intersection of the two laws.

The faster the blanks are filled, the lower the compliance costs will be for business operators. The longer the blanks remain, the longer the situation will persist where business operators must submit different documents to multiple entities for the same act. Erasing the marks of duplication from the Excel file in the meeting room is not the task of lawyers. It is the task of legislation.

Part 6 Supports and Opportunities Available to Business Operators

Chapter 17. Training Data, Infrastructure, and Aggregation Complexes: Laws Granting Government Support

An AI startup trying to build a chest X-ray reading model hit a wall. Although it needed 100,000 medical image data, it received a response that undergoing Institutional Review Board (IRB) review and concluding data use agreements with each individual hospital alone would take eight months. For a company of 12 employees, eight months was enough time to run out of funds. This is why the AI Basic Act codified the integrated provision system for training data into law.

1. Structure of the Integrated Provision System for Training Data

Article 15, Paragraph 4 of the Act stipulates that the Minister of Science and ICT shall establish and manage a system that can comprehensively provide and manage training data—namely, the integrated provision system—and provide it so that the private sector can use it freely. Article 13 of the Enforcement Decree outlines the specific details of this system.

The functions of the system set forth in Article 13, Paragraph 1 of the Enforcement Decree are integrated searching, systematic

classification, and provision of sources for training data, linking with other data platforms operated by state agencies and the private sector, and provision of information related to valuation and quality management. Business operators search for and use necessary datasets in this system.

The Minister of Science and ICT establishes and manages the system. Article 39, Paragraph 2 of the Act provides the legal basis to entrust the establishment, operation, and management of the integrated provision system to a specialized institution. AI Hub (aihub.or.kr), operated by the National Information Society Agency (NIA), has already played a similar role, but Article 15, Paragraph 4 created a direct statutory basis under the name of the integrated provision system.

Use of the integrated provision system is not always free. Pursuant to Article 15, Paragraph 5 of the Act and Article 14 of the Enforcement Decree, the Minister of Science and ICT may collect fees from users and apply differentiated usage fees depending on the type and purpose of use of the training data. If public institutions use it, or if non-profit research institutions or educational institutions use it for academic research or educational purposes, usage fees may be

reduced or exempted. The detailed criteria for charging and targets for fee reduction or exemption are determined by the Minister of Science and ICT through a public notice.

2. Opening of Public Data and Provision for AI Training Purposes

The Public Data Act (Act on Promotion of the Provision and Use of Public Data) has been in effect since 2013. However, this Act was not enacted on the premise of AI training. Public data that is not organized in a standardized format is difficult to feed directly into AI training.

Article 16, Paragraph 2, Item 2-2 of the AI Basic Act enables the state and local governments to provide necessary support to generate and provide data held and managed by public institutions as training data and to secure an appropriate level of quality. Article 13, Paragraph 2 of the Enforcement Decree stipulates that the Minister of Science and ICT may request cooperation from the heads of central administrative agencies, local governments, and public institutions, such as linking their systems and data with the integrated provision system and submitting data, and the heads of the requested agencies shall comply unless there is a special reason. Matters regarding the receipt of public data as training data undergo consultation with the Minister of the Interior and Safety (Article 13, Paragraph 3 of the Enforcement

Decree). Moving beyond the general duty of disclosure under the existing Public Data Act, this opens a path on the government's side for data processed into a form suitable for AI training to flow into the system. The entity making the request is the Minister of Science and ICT, and business operators have not been granted a right to claim processed data from the government.

In practice, this makes a big difference. Take weather data from the Korea Meteorological Administration as an example. The CSV files downloaded from the Public Data Portal contain many missing values, and the regional code system changes from year to year. To use it for AI training, missing value imputation, code integration, and time-series alignment are required, and two engineers may spend three weeks on this preprocessing alone. Receiving datasets processed for training purposes through the integrated provision system reduces these three weeks.

However, there are limitations. Not all public data is immediately provided for AI training. Data involving national security, personal information, or third-party copyrights is excluded. Business operators must first check whether the required datasets are registered in the integrated provision system. If the desired data is not available, they

can raise a request demand with the operating agency, but this is a practical request rather than a legally guaranteed claim procedure.

3. Data Quality Management Support

Article 15, Paragraph 1 of the Act stipulates that the Minister of Science and ICT, in consultation with the heads of relevant central administrative agencies, shall promote policies necessary to secure the quality level of training data, among other things. Article 13, Paragraph 4 of the Enforcement Decree imposes an obligation on the Minister of Science and ICT to make efforts to maintain the timeliness, accuracy, and inter-operability of training data provided through the integrated provision system.

This is a benefit, not an obligation, for business operators. This is because the responsibility for securing quality has been placed on the government's side. However, there is no data quality certification system under the AI Basic Act. Since dataset quality certification is a separate system under the Data Industry Act, business operators needing certification must separately check the procedures under that legal framework.

Quality criteria are important because the performance of AI models directly depends on data quality. A model trained on data containing

labeling errors learns those errors as they are. Improving data quality is often more cost-effective than increasing data quantity.

Pursuant to Article 15, Paragraph 2 of the Act, the government may select target projects for support regarding the production, collection, management, distribution, and utilization of training data and provide support within the scope of the budget. The scale and procedure of support can be confirmed through program announcements published by the Ministry of Science and ICT.

4. Designation Requirements and Occupancy Benefits for Artificial Intelligence Integrated Complexes

Article 23 of the Act establishes the Artificial Intelligence Integrated Complex system. Article 17 of the Enforcement Decree sets forth the designation procedures and considerations.

The concept is as follows: it is a space where AI companies, research institutes, and universities gather in one place to share infrastructure. It can be viewed as an AI-specialized version of Pangyo Techno Valley or Daedeok Innopolis in Daejeon.

Designation is conducted not through a fixed requirement review, but through an evaluation of considerations. Pursuant to Article 17, Paragraph 1 of the Enforcement Decree, the Minister of Science and

ICT or the head of a local government makes designations taking into account matters specified by public notice, such as compliance with the basic plan, effects of regional clustering and competitiveness enhancement of the artificial intelligence industry, potential contribution to local economic development including job creation, redundancy, linkage, accessibility, and efficiency with other integrated complexes, and securing professional personnel and potential for sustainable development. The designation status, location, scope, and boundaries are announced on the website of the relevant institution.

The legal framework for supporting resident companies is Article 23, Paragraph 2 of the Act. The state and local governments may provide administrative, financial, and technical support to designated integrated complexes. Specific benefits, such as rent reductions or provision of computing resources, are not guaranteed by the law itself, but are determined by the support design of each complex. Regulatory exceptions can only be obtained by undergoing procedures under separate legal frameworks, such as the regulatory sandbox, rather than this Act.

As of July 2026, the Artificial Intelligence-Centered Industrial Convergence Integrated Complex Project is operating in Gwangju. For business operators outside the Seoul Metropolitan Area, moving into an integrated complex is a practical route to solving computing infrastructure cost issues.

5. Artificial Intelligence Research Institutes: New Institutions Created by the 2026 Revision

The revised Act of January 20, 2026 (Act No. 21311) newly established Article 22-2: Artificial Intelligence Research Institutes. The same revision also newly established Article 22-3, opening a path for the Minister of Science and ICT to directly establish and operate research institutions to secure innovative artificial intelligence technologies. Article 22-2 entered into force on January 22, 2026, and Article 22-3 entered into force on July 21, 2026.

Artificial intelligence research institutes are not institutions established by the government. They are private research institutes established independently or jointly by universities, enterprises, government-funded research institutes, non-profit corporations, and other entities. They are established with three or more incorporators and authorization from the Minister of Science and ICT, and *mutatis mutandis* apply the provisions governing incorporated foundations in

the Civil Act, except as otherwise provided in this Act. The government and local governments may support operating expenses within the scope of their budgets, and the institutes may receive subsidies and donations or engage in profit-making businesses in accordance with their articles of incorporation.

The statutory projects performed by the institutes are fivefold (Article 22-2, Paragraph 5): research and development of artificial intelligence technology; research and development on convergence across other technologies and disciplines; management, transfer, utilization, and commercialization of research and development achievements; cultivation of professional personnel; and international exchange and cooperation.

What is meaningful to business operators is that the transfer and commercialization of achievements are included among the statutory projects. Even small and medium-sized enterprises lacking resources to build large language models from scratch have a path opened to receive the transfer of research and development achievements from the institutes and apply them to their services. However, specific terms of use are determined through contracts with individual research institutes.

Detailed requirements and procedures for establishment are already provided in the Enforcement Decree. Article 16-2 of the revised Enforcement Decree, which took effect on July 21, 2026, set forth eight types of eligible entities for establishment as well as requirements for financial capability and security measures, while Article 16-3 prescribed the application procedures for establishment authorization. What remains is the Ministry of Science and ICT public notice specifying the granular details of these requirements.

6. Venture Investment Fund of Funds and Promotion of AI Investment

Paragraphs 3 through 5, newly added to Article 18 by the revised Act, are provisions aimed at promoting venture investment in the AI sector. They entered into force on July 21, 2026.

A fund of funds is a large-scale fund formed by combining government contributions and private sector investments. It operates on a two-tier structure where this fund of funds invests in AI-specialized venture capital (VC) funds, which in turn invest in individual AI companies.

Previously, there was already a fund of funds operated by Korea Venture Investment Corp. (KVIC) under the Ministry of Small and Medium Enterprises and Startups (MSS). However, there was no legal

basis specialized for AI utilization. Article 18, Paragraph 3 enables the heads of central administrative agencies, in consultation with the Minister of Small and Medium Enterprises and Startups, to provide support utilizing venture investment funds of funds pursuant to Article 70 of the Venture Investment Promotion Act. Paragraph 4 sets its financial resources as contributions from the state, local governments, and public institutions, contributions from private entities wishing to invest, and incidental revenues. Pursuant to Article 15-3 of the Enforcement Decree, heads of central administrative agencies may request KVIC to formulate an investment plan related to the artificial intelligence industry utilizing the fund of funds. This is not a structure creating entirely new financial resources, but rather one utilizing existing funds of funds for the AI sector.

From the perspective of AI startups, this opens a new path for attracting investment. As the number of AI-specialized VCs backed by fund of funds capital increases, the scale and frequency of investments at the seed round and Series A stages will change. According to figures compiled by the Ministry of Small and Medium Enterprises and Startups, venture investment in the Korean AI sector stood at approximately KRW 1.3 trillion in 2025. Additional government contributions could expand this scale further.

There are practical points that business operators must know. Conditions for fund of funds investment come with restrictions on investment fields and company age/history. In evaluations by AI-specialized VC funds, business operators are often required to present documentation proving that their core technology is AI (such as patents, technical evaluation reports, or records of government R&D projects). Although not a statutory requirement, this is a standard commonly applied in review practices.

7. EU Data Governance Act and China's Data Security Law: Practical Utility of Comparison

South Korea is not the only country establishing a legal basis for AI training data.

The EU's Data Governance Act was enacted and entered into force in 2022 and has applied since September 24, 2023. It is a law regulating conditions for the re-use of public sector data. Although it does not explicitly mention AI training as a purpose, it opens a path to re-use non-personal data for research and innovation purposes within a secure processing environment. It creates a new type of business called data intermediation services, where these intermediaries broker trustworthy transactions between data holders and users.

Korea's integrated provision system is centralized, directly operated by the government. The EU's model is decentralized, establishing private intermediaries managed under a notification and registration system. The underlying philosophies differ.

China's Data Security Law (enacted/implemented in 2021) establishes a data classification and hierarchical protection system. Core state data is managed more strictly, important data is managed through catalogs, and different protection levels are applied depending on the classification grade. The scope of data available for AI training narrows or widens according to this classification. Because cross-border transfers of important data are subject to separate procedures such as security assessments, Korean business operators providing AI services in China often choose to process training locally.

The Interim Measures for the Management of Generative Artificial Intelligence Services (implemented in 2023) require the use of data and foundation models from lawful sources and prohibit infringement on others' intellectual property rights. Upon inspection by regulatory authorities, operators must explain the source, scale, type, and other aspects of training data. Korean law does not yet impose data source explanation obligations to this extent.

If a Korean business operator expands into both markets, it is efficient to design its data governance documentation from the outset to fulfill the requirements of all three jurisdictions.

8. Tension Between Intellectual Property Protection and AI Training

There is no standalone provision for intellectual property protection in the AI Basic Act. The closest provision is Article 13, Paragraph 3. This provision states that when the government discloses or disseminates the results of support projects for the safe use of artificial intelligence, it may protect the entity that developed the technology by setting a protection period to collect technical royalties or through other methods. Intellectual property rights for AI technology itself (patents, trade secrets) and rights in training data are governed by general intellectual property frameworks such as the Patent Act, Unfair Competition Prevention Act, and Copyright Act.

This presents a different angle from the copyright issues addressed in Chapter 14. Copyright prevents unauthorized use of creative works, whereas what is at issue here is the protection of AI technology itself and its output achievements.

To whom do the intellectual property rights of AI models developed by business operators with government support belong? This issue is

not yet clear. Ownership of intellectual property rights for government R&D project achievements follows the Industrial Technology Innovation Promotion Act and National Research and Development Innovation Act, but in the AI field, ownership may differ individually for model weights, training data, and training methodologies.

What happens to the intellectual property rights of derivative models created when business operators fine-tune foundation models developed by Artificial Intelligence Research Institutes (Article 22-2)? Guidelines at the Enforcement Decree level regarding this issue have not yet been issued. When concluding technology transfer contracts with research institutes, business operators must explicitly specify the ownership of rights in derivative models in the contract.

9. What Business Operators Must Prepare to Receive Government Support

The support infrastructure created by the Act broadly consists of four components: the integrated provision system for training data (Article 15), artificial intelligence integrated complexes (Article 23), investment resources utilizing venture investment funds of funds (Article 18, Paragraph 3), and research achievements of artificial intelligence research institutes (Article 22-2). Accessing these four

components requires preparation on the part of business operators.

Applications to use data require a statement of purpose of use, security measures, and personal information processing plans. Operators must specifically describe what they intend to build with the received data and what technical measures will be taken to prevent data leakage.

If data quality certification is required, one must check the quality certification system under the Data Industry Act framework rather than this Act. The AI Basic Act merely places the responsibility for maintaining the quality of data provided through the integrated provision system on the government, without establishing a certification system for business operators' datasets.

Occupancy in integrated complexes follows a public call procedure. Occupancy applications require details on business plans, technical personnel status, and expected performance. Competition rates and evaluation criteria for existing integrated complexes can be confirmed through announcements published by designated institutions (institutions established or designated by the government pursuant to Article 23, Paragraph 4 of the Act) or the relevant local government.

Attracting investment involves checking the list of AI-specialized VCs backed by fund of funds contributions on the KVIC website and submitting a business plan (IR deck) to the relevant VC. Attaching documents proving possession of core AI technologies (such as patent registration certificates, government R&D project completion certificates, or KIBO Technology Fund evaluation reports) will be advantageous in the review process.

Government support is not free. In exchange for receiving support, obligations follow to submit performance reports, undergo post-management, and report data usage performance. Failure to fulfill these obligations constitutes grounds for redemption of support funds. Before applying for support, business operators must factor in the burden of post-management.

Business operators receiving government support domestically while simultaneously expanding into the EU market must check whether government-supported data can be transferred overseas. Some data obtained from the integrated provision system may carry conditions allowing domestic use only. If an EU entity needs direct access to the data, a data transfer agreement should be reviewed; if personal information is included, the operator must satisfy one of the

overseas transfer requirements set forth in the Personal Information Protection Act, such as separate consent, processing entrustment/storage necessary for contract performance, or an adequacy decision.

Chapter 18. Priority Consideration in Public Procurement and Civil Servant Immunity Clause -- The Key to Opening the B2G Market

In the autumn of 2025, an official in charge of the welfare policy division at a ward office in Seoul was unable to answer the ward head's question: "They say an AI system detecting fraudulent receipts of welfare benefits could save billions of won annually, so why haven't we implemented it?" The official's inner thoughts were as follows. If the system malfunctioned and unfairly disqualified legitimate beneficiaries, the official who approved the purchase could be held liable for financial reimbursement in an audit by the Board of Audit and Inspection. The vendor asked to include an immunity clause in the contract, but the legal officer expressed reluctance, citing a lack of precedent. The technology was available and the budget was there, but a single civil servant's personal liability blocked the adoption. The amended AI Basic Act, promulgated on January 20, 2026, added two paragraphs to Article 16 to resolve this deadlock.

1. State Agencies, Consider AI First

Paragraph 3 of Article 16 of the amended Act reads as follows:

"State agencies, etc. shall give priority consideration to artificial intelligence products or artificial intelligence services prescribed by Presidential Decree when intending to purchase products or services or place service contracts necessary for performing their duties: Provided, That this shall not apply where the utilization of artificial intelligence technology is inappropriate due to the nature of the duties."

Though a short sentence, it has significant ripple effects. "State agencies, etc." include the State, local governments, and public institutions under Article 4 of the Act on the Management of Public Institutions (Article 2, subparagraph 4). "Shall" is a mandatory provision. It is not discretionary. It means that when procurement personnel purchase products or services, they must first examine whether an AI alternative exists.

The proviso clause serves as a buffer. In cases where "the utilization of artificial intelligence technology is inappropriate due to the nature of the duties," the obligation is exempted. It does not mean one must consider AI even when purchasing pencils and copy paper.

The limitation "artificial intelligence products or artificial intelligence services prescribed by Presidential Decree" is attached.

Paragraph 4 of Article 15 of the Enforcement Decree defines this as "products or services verified by the Minister of Science and ICT." Not just any AI product qualifies for priority consideration. It must pass through the gateway of verification by the Minister of Science and ICT. The standards and procedures for verification are publicly notified by the Minister of Science and ICT in consultation with the Minister of the Interior and Safety (Article 15, Paragraph 5 of the Enforcement Decree).

The proposed amendment to the Enforcement Decree, announced for advance notice of legislation on May 21, 2026, specified this verification procedure. Under this structure, the Korea Artificial Intelligence Industry Association verifies whether AI technology is applied to the product, and the Telecommunications Technology Association (TTA) conducts technical reviews. Detailed procedures are expected to be determined according to the standards finalized and publicly notified by the Minister of Science and ICT.

Unpacking this structure from a business perspective leads to this: to supply an AI product to the public market, it must undergo reviews by the Korea Artificial Intelligence Industry Association and TTA, and obtain verification from the Minister of Science and ICT. A verified

product is examined ahead of non-AI alternatives during the procurement process of State agencies.

Similar systems previously existed in Korea. These include the public purchase support system for innovative products under Article 27 of the Government Procurement Act (enforced in 2020), the Public Procurement Service's innovative product pilot purchase project, the Innovation Market (the comprehensive information system for public procurement of innovative products), and the Ministry of Economy and Finance's directive Regulations on the Designation and Purchase Promotion of Innovative Products. The existing innovative product system granted discretion to the Administrator of the Public Procurement Service, stating that they "may support" the public purchase of innovative products. Paragraph 3 of Article 16 of the AI Basic Act imposes a mandatory requirement: "shall give priority consideration." It has shifted from discretion to an obligation.

2. Immunity as a Safety Device

Let us examine Paragraph 4 of Article 16.

"Even if loss or damage occurs to a State agency, etc. due to the purchase or use of an artificial intelligence product or artificial intelligence service under Paragraph 3, the person in charge of such

purchase or use duties shall not be liable to indemnify the loss or damage incurred by the relevant agency: Provided, That this shall not apply where loss or damage occurs due to intention or gross negligence of the person in charge of the relevant duties."

To understand why this provision was necessary, one must consider the position of civil servants. AI operates probabilistically. Outcomes vary depending on training data, and perfect verification is impossible at the time of adoption. Traditional software yields the exact same output for a given input. AI is different. Different outputs based on probability distributions can result from the same input.

What happens if an agency suffers damages due to a malfunction after a civil servant adopts an AI system? The appropriateness of budget execution under the National Finance Act could be called into question. If an audit by the Board of Audit and Inspection points out that "it was adopted without sufficient review," the person in charge along the approval chain becomes subject to accountability. Disciplinary action under the Decree on Disciplinary Action against Public Officials and the exercise of the right to indemnity under the State Compensation Act are also possible.

What would a rational civil servant do? Maintain the existing system. Not introducing new technology is personally safer. In public administration, this is called 'risk-aversion bias.' It is a common cause of slow innovation in the public sector. Doing nothing incurs no liability. Trying something and failing incurs liability. In this asymmetric structure, innovation grinds to a halt.

Paragraph 4 of Article 16 corrects this asymmetry by law. A civil servant who purchases or uses an AI product verified by the Minister of Science and ICT will not be held liable for indemnification even if damage is incurred by the agency as a result.

The conditions for immunity are twofold. First, the purchased product must qualify for priority consideration under Paragraph 3. This means it is a product verified by the Minister of Science and ICT. If an unverified AI product is arbitrarily purchased, it is not eligible for immunity. Second, there must be no intention or gross negligence.

Let us draw the boundary line specifically.

Suppose an AI recruitment evaluation system was adopted, resulting in outcomes disadvantageous to a specific gender. The manufacturer recommended "conducting quarterly bias checks," but

the official did not perform a single check for two years. This could constitute gross negligence.

Conversely, if the same system was adopted and quarterly checks were conducted according to the manufacturer's manual, but an error occurred due to an unexpected data anomaly, immunity applies.

A precedent with a similar structure exists. Paragraph 4 of Article 27 of the Government Procurement Act provides that the purchasing manager of a demanding organization who purchases an innovative product shall not be held liable for losses incurred by purchasing the product unless intention or gross negligence is proven. The structure—granting immunity to purchasing managers while making intention and gross negligence exceptions—is identical. Paragraph 4 of Article 16 of the AI Basic Act stands upon this lineage while directly addressing the new problem of AI's probabilistic uncertainty.

This holds significance for businesses. A major psychological barrier that made public clients hesitant to purchase AI has been lowered by law. Sales representatives can now explain in the field: "As long as there is no intention or gross negligence, you will not have to compensate for agency losses; the law contains an immunity clause."

3. The Second Door Opened by Impact Assessment

The priority consideration obligation under Article 16 applies to AI products and services verified by the Minister of Science and ICT. For high-impact AI, there are two additional priority consideration mechanisms.

Paragraph 2 of Article 35 states: "Where State agencies, etc. intend to use products or services utilizing high-impact artificial intelligence, they shall give priority consideration to products or services for which an impact assessment has been conducted."

The impact assessment itself is not mandatory. Paragraph 1 of Article 35 stipulates that they "shall endeavor to." This 'duty to endeavor' is an obligation carrying no sanctions. Even if not fulfilled, no administrative fine is imposed. In fact, examining Article 43 on administrative fines reveals that administrative fines apply to only three categories: violation of transparency prior notification obligations (Article 31, Paragraph 1), failure to designate a domestic representative (Article 36, Paragraph 1), and non-compliance with suspension orders or corrective orders (Article 40, Paragraph 3). There is no provision imposing an administrative fine for failing to conduct an impact assessment.

However, Paragraph 2 adds market weight to this autonomy. Products that have completed an impact assessment receive prior consideration in public procurement. Products that have not are pushed behind.

Article 28 of the Enforcement Decree lists seven items that must be included in an impact assessment: identification of individuals or groups likely to be affected (including reflection of vulnerable group characteristics), identification of fundamental rights categories, the content and scope of social and economic impacts, usage behavior, quantitative or qualitative evaluation metrics and output generation methods, measures for risk prevention, mitigation, and loss recovery, and an implementation plan for improvement. A business operator may conduct it directly or commission a third party (Article 28, Paragraph 2 of the Enforcement Decree).

Paragraph 4 of Article 30 employs the same logic: "Where State agencies, etc. intend to use high-impact artificial intelligence, they shall give priority consideration to products or services based on artificial intelligence that has received verification, certification, etc." AI products that have undergone voluntary verification or certification regarding safety and reliability receive additional priority

consideration in public procurement.

Placing these three provisions side by side reveals the structure.

Paragraph 3 of Article 16 prescribes priority consideration for AI products verified by the Minister of Science and ICT.

Paragraph 4 of Article 30 prescribes priority consideration for high-impact AI products that have received verification and certification.

Paragraph 2 of Article 35 prescribes priority consideration for high-impact AI products for which an impact assessment has been conducted.

It is a triple layer. A product that satisfies all three—Ministry of Science and ICT verification, verification/certification, and impact assessment—passes through three gateways simultaneously. A product that lacks even one faces disadvantageous competitive conditions in the public market.

4. The Logic of Soft Regulation

This is the 'soft regulation' structure of the AI Basic Act.

The statute does not directly mandate impact assessments or certifications. It does not impose administrative fines. Instead, it induces voluntary compliance through the medium of large-scale demand in the government procurement market. There is no punishment for non-compliance with the law, but complying with the law wins the market.

The ceiling for administrative fines under Korean law is 30 million won (Article 43). The cap on fines under the EU AI Act varies depending on the type of violation. Violating prohibited AI practices (Article 5) incurs fines up to 35 million EUR (approximately 50 billion KRW) or 7% of total worldwide annual turnover, while violating business operator obligations such as launching high-risk AI without a conformity assessment incurs up to 15 million EUR or 3% of total worldwide annual turnover. Even between Korea's 30 million won and the EU's 15 million EUR, there is roughly an 800-fold gap.

Korean legislators intentionally created this disparity. Based on the judgment that the domestic AI industry is in an early stage of growth, they accepted the corporate argument that excessive sanctions could stifle innovation. While the whip is small, the carrot is large. An administrative fine of 30 million won is unlikely to motivate large

corporations to comply, but the prospect of losing the entire government procurement market serves as a compelling incentive for compliance.

This approach is not unique to Korea. Using government procurement as an incentive for regulatory compliance also appeared in the United States. The Biden administration's Executive Order 14110 on AI in 2023 directed federal agencies to give priority consideration to products meeting safety and transparency requirements when procuring AI (this Executive Order was revoked in January 2025). Article 16 of Korea's AI Basic Act codifies a similar logic at the statutory level.

There is a difference. In the United States, because it was an Executive Order, it could be discarded when the presidency changed. In January 2025, the Trump administration revoked Biden's AI Executive Order. Korea's framework is statutory law. It remains in effect unless amended by the National Assembly.

5. Scale of the B2G Market

The AI demand in Korea's public sector can be discerned through institutional indicators.

Chief Artificial Intelligence Officers (CAIOs) have been designated across all 48 central administrative agencies. Ministerial-level agencies appointed vice-ministerial officials, while vice-ministerial organizations such as administrations appointed director-general level officials. They have held monthly regular consultative meetings starting September 18, 2025. The National AI Strategy Commission operates under a framework chaired by the President, with a maximum of 60 regular members, 52 subcommittee members across 8 subcommittees, totaling 85 members.

On February 25, 2026, during the 2nd plenary meeting of the Strategy Commission, the 'Republic of Korea AI Action Plan (2026–2028)' was finalized. It comprises 99 execution tasks and 326 policy recommendations. Heads of the 48 agencies will execute the 99 tasks within three years. A significant portion of these tasks is highly likely to translate into the purchase of AI products and services.

The AI Policy Center was established within the National Information Society Agency (NIA) on April 29, 2026, and officially opened on June 30. It operates with a staff of 59 personnel across 3 divisions, 1 group, and 5 teams. Its statutory missions include

supporting basic plan formulation, supporting policy development, and conducting surveys on the impact of AI dissemination. The AI Safety Institute (AISi) operates separately as an affiliate of the Electronics and Telecommunications Research Institute (ETRI), taking charge of safety evaluations and safety technology research. Policy planning institutions, safety research institutions, and business operator support institutions are all in place. All that remains is actual procurement.

The priority consideration obligation under Article 16, Paragraph 3 takes effect on July 21, 2026. From this day forward, procurement personnel at central administrative agencies, local governments, and public institutions must first examine AI alternatives when purchasing products or services or contracting services. Combining Korea's State agencies, local governments, and public institutions yields thousands of entities. Demand for AI transformation in each institution's IT projects, administrative services, and citizen-facing operations is created by law.

Korea is a country that has consistently ranked near the top of the UN E-Government Development Index (EGDI) for years. Digital public service infrastructure such as Government24, National Assistant, and

National Tax Service HomeTax is already established. Layering AI on top of this infrastructure is technically a natural next step. Article 16, Paragraph 3 has placed a statutory obligation on this next step.

6. What Businesses Must Prepare

Here is a step-by-step summary of what businesses seeking to supply AI to the public sector must verify.

First, determine whether your product qualifies as high-impact AI. Article 2, subparagraph 4 of the Act enumerates 11 domains (energy supply, healthcare, medical devices, nuclear safety, biometric information analysis, determinations significantly affecting rights such as recruitment or loan screening, traffic management, qualification screening and service fee collection by public institutions, student evaluation, etc.). If it falls into these domains, it is high-impact AI. If applicability is unclear, a request for confirmation may be submitted to the Minister of Science and ICT (Article 33, Paragraph 1).

If it is a general AI product that does not qualify as high-impact AI, obtaining only the verification of the Minister of Science and ICT (Article 15, Paragraph 4 of the Enforcement Decree) makes it eligible for priority consideration. This is a procedure involving technology

application verification by the Korea Artificial Intelligence Industry Association and a technical review by TTA.

If the product qualifies as high-impact AI, it is advantageous to satisfy all three priority consideration provisions.

Ministry of Science and ICT verification (Article 15, Paragraph 4 of the Enforcement Decree in connection with Article 16, Paragraph 3) is fundamental.

Conducting an impact assessment (Article 35) includes the product in the second priority consideration category. An assessment report covering the 7 items specified in Article 28 of the Enforcement Decree must be prepared. It must identify individuals or groups likely to be affected (reflecting vulnerable group characteristics), identify fundamental rights categories, analyze the scope of social and economic impacts, describe usage behavior, present evaluation metrics and output generation methods, prepare measures for risk prevention, mitigation, and loss recovery, and include an implementation plan for improvement. Businesses may perform this directly or commission a third-party specialized institution.

Obtaining verification and certification (Article 30) includes the product in the third priority consideration category. This refers to voluntary safety and reliability verification or certification supported by the Minister of Science and ICT.

Fulfilling all three satisfies all three priority consideration provisions. A competitive edge in public procurement becomes clear.

There is one more item to check: the AI Basic Act Support Desk. Operated by the Korea Software Industry Association (KOSA), it received 552 consultations as of March 31, 2026, the 10th week of implementation. Anonymous consulting is available. Trade secrets are protected. Businesses can receive advice on whether a product qualifies as high-impact AI or on methods for conducting impact assessments.

7. The EU Chose Sandboxes

Instead of priority consideration in public procurement, the EU AI Act introduced a mechanism known as a regulatory sandbox. Articles 57 through 62 govern this system.

A regulatory sandbox is a controlled environment where business operators can develop and test AI systems under regulatory

supervision before introducing them to the actual market. The EU AI Act mandated that each member state establish at least one AI regulatory sandbox. The original deadline was August 2, 2026, but it was postponed by one year to August 2, 2027, under the Digital Omnibus amendment.

Spain operating a pilot sandbox in June 2022 even before the legislation was finalized became the EU's first case. However, as of July 2026, the majority of member states have not yet officially opened an AI regulatory sandbox. The one-year extension of the deadline itself reveals a lack of preparation.

Inside the sandbox, enterprises test high-risk AI systems under limited conditions while receiving supervision and guidance from regulatory authorities. Preferential measures, such as fee reductions, are available for SMEs and startups. Real-world testing under restricted conditions is also permitted, in which case prior consent from users is required (Articles 60–61).

The approaches of Korea and the EU differ.

The EU takes an ex-ante regulatory approach. Before high-risk AI can be placed on the market, it must undergo a conformity

assessment. The regulatory sandbox is a laboratory for preparing for this conformity assessment. Regulation comes first, and the sandbox is a tool to pass regulation.

Korea takes a market-incentive approach. It does not mandate impact assessments or certifications by law. Instead, it induces voluntary compliance through the medium of the public procurement market. Regulation stands behind, and the market stands in front.

Korea also has a foundation for practical demonstration. Article 24 of the Act stipulates the "creation of a foundation for artificial intelligence demonstration," enabling the State and local governments to construct and operate facilities, equipment, and installations to support demonstration and performance testing. While outwardly similar to the EU's regulatory sandbox, its purpose is different. The EU sandbox is an environment designed to help with regulatory compliance. Korea's demonstration foundation is an infrastructure for industrial fostering.

There are also critiques of the EU sandbox logic. There is no guarantee that an AI system passing inside a sandbox will remain safe in the complex environment of the real market. Success in a

controlled environment does not translate to success in an uncontrolled environment. On the other hand, there are critiques of Korea's soft regulation as well. Companies with no intention of entering the public procurement market have no incentive to conduct impact assessments. The carrot of soft regulation does not reach AI products exclusively meant for the private market. Under Korean law, an administrative fine of 30 million won is an absurdly small amount to enforce compliance in the private domain.

8. The Door Has Opened

After July 21, 2026, Korea's public sector bears a legal obligation to give priority consideration to verified AI products. Civil servants receive the protection of immunity. Businesses must obtain the qualification of verification from the Minister of Science and ICT.

The welfare policy official at that ward office in Seoul now has a legal answer. They can introduce an AI welfare fraud detection system, and even if loss occurs due to malfunction, as long as it was operated according to the manual, the official will not be held liable to indemnify agency losses. There is a condition: it must be a product verified by the Minister of Science and ICT.

The law has opened the door. Who will pass through that door depends on the verification standards of the Enforcement Decree notification, the review capacity of the Korea Artificial Intelligence Industry Association and TTA, and the quality of products brought forth by businesses. The public market is slower than the private market, has lengthy procedures, and requires familiarity with administrative language. Instead, once entered, contracts span multiple years, and the speed of dissemination across institutions is rapid. An AI service that succeeds in one ministry can spread to other ministries through the CAIO Consultative Council across 48 agencies.

What the law has created is the door to the market. The market itself is created by products.

Chapter 19. Accessibility Support for Vulnerable Groups and Cost Subsidies: A New Market Opened by the 2026 Amendment

Let us picture 78-year-old Grandmother Kim, living in a rural village in South Jeolla Province, attempting to use an agricultural AI pest and disease diagnosis app. She opened the app installed by her granddaughter, but the font was too small, and the voice recognition failed to properly understand the local dialect. After three attempts, she gave up. No matter how sophisticated AI technology may be, for those who cannot access it, it is as good as nonexistent. This is why the 2026 amended Act newly established Article 17-2 to pave the way for supporting usage costs.

1. An Article Absent Prior to the Amendment: Article 17-2 Newly Created by the Revision

In the original Act (promulgated on January 21, 2025, and entering into force on January 22, 2026), there was no independent article dealing with AI accessibility or cost support for vulnerable groups. The extent of its provisions was merely requiring the inclusion of accessibility assurance within the basic plan.

The amended Act of January 20, 2026 (Act No. 21311, entering into force on July 21, 2026) filled this void. Article 3, Paragraph 5 refers to

persons with disabilities, elderly persons, and others who experience difficulties in using artificial intelligence products or artificial intelligence services as "vulnerable groups to artificial intelligence," imposing a duty on the State and local governments to guarantee their participation and strive to reflect their opinions in the process of developing and establishing artificial intelligence policies. Article 6, Paragraph 2, Item 7 made it mandatory for the basic artificial intelligence plan to include matters for guaranteeing the access and use of vulnerable groups. Furthermore, Article 17-2 created a new legal basis allowing the State to subsidize all or part of the costs, within the scope of the budget, for persons who have difficulty using AI products and services due to economic conditions.

The scope of vulnerable groups is defined into ten categories under Article 1-2 of the Enforcement Decree: persons eligible for benefits under the Act on the Promotion of Employment and Livelihood Stability of Jobseekers; beneficiaries of basic livelihood security and the near-poverty class; farmers and fishermen; members of multicultural families; North Korean defectors; women with career background; persons with disabilities; persons eligible for support under the Single-Parent Family Support Act; persons aged 65 or older; and persons recognized by the head of a central administrative

agency or a local government as experiencing difficulties in usage. The eligibility criteria for cost support are separately prescribed under Article 15-2 of the Enforcement Decree, which encompasses science and engineering talent, persons of national merit rated with a disability grade, and regional balanced talent in addition to the aforementioned categories.

What was once a mere declaration has now transformed into a legal basis for cost support. What this change signifies for business operators is clear: the government can now allocate budget to subsidize AI service usage fees for specific target groups. From the perspective of business operators, a new demand base funded by government subsidies has emerged.

2. The Size of the Market Opened by Cost Support

Let us look at the numbers. Registered persons with disabilities account for approximately 2.63 million, and elderly persons aged 65 or older account for approximately 10.5 million. As of the end of 2024, beneficiaries of basic livelihood security number about 2.67 million; career-interrupted women—the core group of women with career background—stand at around 1.1 million as of 2025; and farmers and fishermen are estimated at around 2.3 million. Added to these are members of multicultural families, North Korean defectors, persons

eligible for single-parent family support, and the near-poverty class, while science and engineering talent, national merit recipients with disability ratings, and regional balanced talent are included as targets for cost support.

The term "jobseekers" must be interpreted with caution. The target designated by the Enforcement Decree is not jobseekers in general, but rather persons eligible for benefits under the Act on the Promotion of Employment and Livelihood Stability of Jobseekers—that is, individuals recognized as eligible beneficiaries under the National Employment Support Program. Because this is a distinct system from beneficiaries of unemployment benefits under employment insurance, the scales of the two groups should not be commingled in calculations.

Even after eliminating overlaps, the potential beneficiary population reaches millions of people. If one assumes that the government issues AI service vouchers worth 10,000 KRW per person per month, a market worth hundreds of billions of KRW per year would be created. However, this calculation rests entirely on assumptions. The law merely established the legal ground that support "may be provided," leaving the adoption of vouchers, the

specific amounts, and the scale of beneficiaries completely undecided.

The actual size will depend on budget allocations. Immediately following implementation in the second half of 2026, it is highly likely that a small budget will be assigned in the form of pilot projects. Nevertheless, the policy trajectory points toward expansion. Once support for AI service usage comes into full swing, it will serve as a stable demand foundation for AI business operators in the education, healthcare, and agriculture sectors.

3. AI Accessibility for Persons with Disabilities: What to Guarantee and How

The AI Basic Act and its Enforcement Decree have not yet separately prescribed technical standards for accessibility for persons with disabilities. Standards prevailing in practice derive from the Act on the Prohibition of Discrimination against Persons with Disabilities, Remedy for Rights Infringement, etc., its Enforcement Decree, and the Korean Web Content Accessibility Guidelines (KWCAAG). Applying these standards to AI services reveals what must be prepared.

For visually impaired individuals, an interface compatible with screen readers is required. If an AI chatbot responds in text, that text must be readable by a text-to-speech (TTS) engine. For generative AI

that creates images, a function to attach alternative text (alt text) to generated images is necessary. Because automatically generated alt text may lack accuracy, it is safer to incorporate a human review process.

For hearing-impaired individuals, voice-based AI services require subtitles or sign language interpretation features. If an AI voice assistant provides responses exclusively through audio without a text-conversion option, hearing-impaired users cannot utilize that service.

For individuals with physical disabilities, an interface navigable solely by keyboard or controllable via voice commands is necessary. An AI service where mouse clicks or touchscreen operations are the sole input methods fails to meet accessibility standards.

These requirements are not novel; rather, they represent an expansion into the AI domain of what existing anti-discrimination laws for persons with disabilities and web accessibility standards have long mandated. However, given the nature of AI services, additional considerations are warranted. Who is responsible for the accessibility of content produced by generative AI? If a user business operator uses an AI-generated marketing image that lacks alt text,

does this constitute a breach of accessibility obligations by the user business operator, or a functional deficiency on the part of the generative AI developer business operator? This boundary of responsibility has not yet been clearly settled anywhere in the legislation.

4. Disseminating AI to Senior Citizens: Design Overcoming the Digital Divide

For AI services targeting senior citizens, interface design—rather than the technology itself—is the key determinant.

According to the Survey on the Digital Divide conducted by the Ministry of Science and ICT and the National Information Society Agency, smartphone ownership among elderly persons in their 70s and older has risen to around 90%. However, when setting the digital capability of the general public at 100, their utilization level remains in the 40s. They possess devices but cannot utilize them effectively.

For an AI service business operator to enter the senior market, large fonts, high contrast, and voice-first interfaces are fundamental requisites. In addition, these must be backed by design enabling easy reversal of mistakes (clear placement of an undo function), instruction text free of technical jargon, and human customer support representatives accessible via a single phone call.

The inclusion of persons aged 65 or older under Item 9 of Article 1-2 of the Enforcement Decree as a vulnerable group to artificial intelligence serves as a signal that government support will be linked to business operators implementing such designs. Business operators developing AI services tailored for the elderly would do well to examine business announcements regarding digital inclusion issued by the Ministry of Science and ICT to check eligibility requirements for participation.

5. Women with Career Background and Jobseekers: AI Education and Service Support

The inclusion of women with career background and eligible jobseeker beneficiaries among vulnerable groups is unusual. This is because the classification is based on economic and social circumstances rather than physical conditions such as disability or age. The statutory terminology is also noteworthy: rather than the commonly referenced "career-interrupted women," the law employs the term "women with career background, etc." under the Act on Prevention of Career Interruption and Promotion of Economic Activities of Women, while jobseekers are restricted to eligible benefit recipients under the National Employment Support Program.

The legislative intent is to prepare for structural employment shifts brought about by AI. If the transition speed of the workforce capable of utilizing AI lags behind the pace at which AI replaces existing jobs, individuals will be marginalized by that gap. Women with career background and jobseekers represent groups highly susceptible to falling into this gap.

For business operators, this signifies an expansion of the AI education service market. If the government issues AI utilization training vouchers to these groups, business operators will be needed to accept those vouchers and provide training. Demand will emerge for prompt engineering training, instruction on using AI assistance tools, and hands-on practice in applying AI across various industries.

If linked with the Ministry of Gender Equality and Family's support projects for women with career background and the Ministry of Employment and Labor's National Tomorrow Learning Card system, the budgetary scale could grow substantially. However, such linkage is not a statutory effect prescribed by law, but a possibility contingent upon policy design. Business operators must verify whether their AI education programs can be incorporated into this government support framework. Accreditation requirements follow the

educational institution registration criteria of the relevant government ministries.

6. Farmers, Fishermen, and Regional Balanced Talent: On-Site Dissemination of AI

The inclusion of farmers and fishermen as target beneficiaries (Item 3 of Article 1-2 of the Enforcement Decree) reflects the policy goal of mitigating regional disparities.

Small and medium-sized enterprises (SMEs) outside the capital region lack personnel and infrastructure even when willing to adopt AI. Commissioning an AI consulting firm based in Seoul costs millions of KRW in travel expenses alone. While cloud GPU costs remain identical to those in the capital region, talent capable of refining data and deploying models on-site is absent locally. However, a support category explicitly named "employees of SMEs outside the capital region" does not exist in the legislation. What the Enforcement Decree designates is regional balanced talent referring to graduates from regional universities (Item 4 of Article 15-2) and persons recognized by the head of a central administrative agency or local government as experiencing usage difficulties (Item 10 of Article 1-2, and Item 5 of Article 15-2). Although local governments could potentially encompass regional industrial personnel under this

discretionary recognition clause, that remains subject to confirmation once official announcements are published.

The circumstances facing farmers and fishermen are even more severe. While AI application fields such as smart farming, pest and disease diagnosis, crop yield prediction, and aquaculture water quality management are clear, local digital infrastructure and user capability lag behind.

Government cost subsidies narrow this gap. What Article 17-2 directly establishes is support for usage costs. On-site tailored training or device distribution can only be attached if supported by separate project rationales and budgets. Nevertheless, it is readily foreseeable that multiple forms of support may be bundled into a single announcement, much like local government smart farm support programs. For smart farm AI business operators, agricultural drone analysis business operators, and aquaculture management AI business operators, a stable customer base backed by government subsidies is established.

To enter this market, business operators must make several preparations: track relevant support program announcements from the Rural Development Administration, the Ministry of Oceans and

Fisheries, and the Ministry of Small and Medium Enterprises and Startups; apply to participate as AI service providers in smart farming and smart aquaculture support projects launched by local governments outside the capital region; and equip services with offline modes or low-bandwidth environment compatibility for practical on-site operation.

7. Comparison with the EU Accessibility Act and China's Barrier-Free Environment Construction Law

The European Accessibility Act (Directive (EU) 2019/882) became applicable on June 28, 2025. This law does not target AI exclusively; it regulates the accessibility of enumerated products and services such as computers, smartphones, ATMs, e-commerce websites, and transportation services. If an AI service is integrated into these listed products or services, it applies jointly, meaning Korean AI business operators entering the EU market must satisfy accessibility requirements within that scope.

Under the EU AI Act, accessibility of high-risk AI is addressed in Article 16. Providers of high-risk AI systems bear an obligation to design systems to satisfy accessibility requirements pursuant to the Public Sector Web Accessibility Directive (Directive (EU) 2016/2102) and the European Accessibility Act (Article 16, Point (I)).

The distinction between Korean law and EU law lies in their approach. Korea employs demand-side support that subsidizes costs for vulnerable groups. The EU enforces supply-side regulation mandating business operators to satisfy accessibility standards. If a Korean business operator enters the EU market, it faces a dual structure: acting as a beneficiary of demand generated by cost subsidies domestically, while being an addressee of accessibility obligations in the EU.

China's Law on the Construction of a Barrier-Free Environment (effective September 2023) guarantees information accessibility for persons with disabilities and the elderly. It mandates barrier-free transformation of websites and mobile applications; although AI services are not explicitly specified, as long as AI is provided via websites and applications, there is strong room for interpretation that identical demands extend to them. Sanctions center on corrective orders and public criticism notifications, with administrative fines restricted to breaches of duty by certain entities such as telecommunications operators.

Among the three jurisdictions, Korea's Article 17-2 is unique in terms of target scope. While accessibility legislation in the EU and China

focuses on persons with disabilities and the elderly, Korea encompasses socioeconomically vulnerable groups such as women with career background, eligible jobseeker beneficiaries, farmers and fishermen, and North Korean defectors. Whether this expansion will be backed by actual budgets and policies remains a key observation point following the Act's enforcement.

8. Size of the AI Accessibility Market and Entry Strategies

Determining the precise market size at this stage is difficult. Government budgets have not been finalized, and support targets and criteria are established by heads of central administrative agencies and local governments through official announcements pursuant to Article 15-2 of the Enforcement Decree. Announcements will contain criteria reflecting applicants' necessity for AI utilization, economic, regional, physical, and social circumstances, budget scale, and equity in support distribution.

Nonetheless, the direction can be discerned. The Ministry of Science and ICT has annually operated digital inclusion programs. If support for AI usage costs is integrated into this program framework, it is likely to launch at a pilot project scale initially, with potential for expansion once outcomes are verified.

Entry strategies for business operators vary by target demographic. For persons with disabilities and senior citizens, technological capabilities such as assistive device integration, voice interfaces, and automatic alt-text generation are paramount. For women with career background and jobseekers, educational content quality and post-completion employment placement outcomes serve as evaluation metrics. For farmers and fishermen, practical field applicability (resilience in unstable internet environments, waterproof/dustproof device compatibility) is decisive.

Certain common preparations should be made: familiarize oneself in advance with competition and public announcement procedures for government support projects; obtain accessibility certification for persons with disabilities (Web Accessibility Quality Mark)—which, though not a legally mandatory certification, serves as evidence of a prepared business operator in public and welfare tenders; and establish a user satisfaction measurement mechanism. Because government-supported projects are frequently evaluated on performance metrics such as usage counts, satisfaction levels, and reuse rates, a system to collect these metrics should be embedded into the service.

9. Relationship with Existing Accessibility Legislation

Article 17-2 does not operate in isolation. It interconnects with existing legislation, including the Act on the Prohibition of Discrimination against Persons with Disabilities, Remedy for Rights Infringement, etc., the Act on Welfare of Persons with Disabilities, the National Basic Livelihood Security Act, the Multicultural Families Support Act, the North Korean Refugees Protection and Settlement Support Act, the Single-Parent Family Support Act, the Act on Prevention of Career Interruption and Promotion of Economic Activities of Women, and the Special Act on the Improvement of Quality of Life of Farmers and Fishermen.

The Act on the Prohibition of Discrimination against Persons with Disabilities already prohibits discrimination in information access. If an AI service is provided in a form inaccessible to persons with disabilities, potential violations of this Act may arise depending on the service type and the scope of duty to provide reasonable accommodation. Article 17-2 is not a prohibitory norm, but rather an article adding affirmative support (subsidies for usage costs) on top of it.

The law does not establish a priority of precedence between the two statutory systems. Nevertheless, the practical sequence is evident.

Seeking to participate in government support projects while subject to recommendations for corrective action due to accessibility deficiencies under anti-discrimination laws lacks persuasive standing. Fulfilling the legal obligation of non-discrimination first, and subsequently seizing government support opportunities (cost subsidies, program participation) on that foundation, represents the safe path.

The Special Act on the Improvement of Quality of Life of Farmers and Fishermen includes the expansion of information and communications infrastructure in rural and fishing communities as a public service item. For Article 17-2 cost support to reach farmers and fishermen, budget and project coordination between the Ministry of Agriculture, Food and Rural Affairs and the Ministry of Science and ICT serves as a practical variable.

Accessibility for vulnerable groups is not a burden for business operators, but an opportunity. This is because it presents a structure where the government creates demand, subsidizes costs, and entrusts supply to business operators. Whether this framework operates effectively will be determined by public announcements from respective government ministries and local governments, as

well as budgetary allocations.

"□□□ □□□□ □□□ □□□□ □□, □□□ □□ □□. □□ □□□
□□□ □□ □□ • □□ • □□□"

Chapter 20. Startups, Specialized Personnel, Venture Investment Fund of Funds, and AI Research Institutes -- How to Secure Capital and Talent

In the autumn of 2025, three CEOs of AI startups had lunch together at a shared office in Gangnam, Seoul. One was developing medical image diagnostic AI, another was developing a manufacturing process optimization model, and the third was preparing an educational generative AI. Their shared concern was not technology; it was money and talent. Seed investment funds were running out, yet they could not find next-round investors, and hiring master's-level AI engineers required competing on salary with major conglomerates. Although they were aware that the AI Basic Act would take effect on January 22, 2026, they only thought of regulatory obligations and never imagined that the law could offer them anything.

The AI Basic Act is not merely a regulatory law. Out of the 46 articles in the entire act, 17 articles from Article 13 to Article 26 are dedicated to technology development and industrial promotion. The amended act of January 20, 2026 added four additional grounds: utilization of venture investment funds of funds, support for usage costs for vulnerable groups, provision of public data as training data, and research institutions established directly by the government. The

legal grounds regarding the capital and talent that the three CEOs were searching for had already been established.

1. Startup Support: Seven Pathways Opened by Law

Article 18 of the AI Basic Act lists seven types of projects under the title "Activation of Startups, etc." These are: discovery and nurturing of founders, startup education and training, support for the commercialization of superior technology by specialized personnel, valuation of AI technology and financial support for startup capital, provision of research and technology development outcomes, fostering of startup support institutions and organizations, and other projects necessary for activating startups.

There are two points to pay attention to here. One is "valuation and financial support of startup capital" in Article 18, Paragraph 1, Item 4. A hurdle that AI technology startups find difficult to overcome is proving the value of their technology in numbers. They have almost no tangible assets, no revenue yet, and only model performance metrics and datasets. In effect, the law has provided the government with the legal basis to establish a system for evaluating the value of AI technology.

The other is Article 18, Paragraph 2. It stipulates that local governments may make contributions or equity investments in public institutions or public organizations that support AI startups. This means that the Seoul Metropolitan Government or Gyeonggi-do can establish accelerator institutions dedicated to AI startups and allocate budget to them. Including basic local governments, a legal basis has been created for public funds to be injected into AI startup support institutions anywhere across the 243 local governments nationwide.

The contrast is vivid. The EU AI Act is a regulatory law. Although there are consideration measures for small and medium-sized enterprises (SMEs) such as priority access to regulatory sandboxes, there are no provisions directly supporting startup capital or investment. The EU promotes AI industrial development through a separate policy package called InvestAI. Announced by the European Commission in February 2025, this package targets a combined public-private investment of 200 billion euros. South Korea, which included promotion provisions within the law, and the EU, which separated regulatory law from promotion policy, differ fundamentally in their approaches.

China is different again. The State Council's "Next Generation Artificial Intelligence Development Plan" of 2017 forms the root of AI industrial promotion. The starting point is a State Council document rather than a statute, operating under a structure where local governments competitively provide industrial subsidies, land, and tax cuts. That the city of Shenzhen provides research subsidies to AI companies and the city of Hangzhou gives housing subsidies to AI talent stems not from statutory mandates, but from discretionary industrial policies of local governments. Article 18 of South Korea's AI Basic Act differs in structure in that it has secured the legal basis for such support at the level of a basic act.

2. Venture Investment Fund of Funds: The Structure Through Which Government Money Reaches Startups

Article 18, Paragraph 3 is the core: "The head of a central administrative agency may, in consultation with the Minister of SMEs and Startups, provide support utilizing venture investment funds of funds in order to activate startups in the field of the artificial intelligence industry."

A venture investment fund of funds is a fund of funds based on Article 70 of the Venture Investment Promotion Act. It is managed by Korea Venture Investment Corp. The government and public

institutions contribute capital to the fund of funds, which then places those funds into sub-funds such as venture investment partnerships or individual investment partnerships. The sub-funds invest in individual startups. It is a structure where government money does not go directly to individual companies, but flows through professional general partners (GPs).

The AI Basic Act has connected this existing structure so that it can be used exclusively for AI. Article 15-3 of the Enforcement Decree stipulates concrete procedures. Reflecting the outcome of consultations with the Minister of SMEs and Startups, the head of a central administrative agency may request Korea Venture Investment Corp. to establish an investment plan related to the AI industry. This means that the Minister of Science and ICT can be the requesting party.

The composition of financial resources is stated in Article 18, Paragraph 4. Capital contributions from the state, local governments, and public institutions form the foundation, to which money from private contributors and incidental revenues are added. It is a design aimed at a leverage effect where the government injects seed money and private capital follows by multiples.

What does this mean from the perspective of a startup? If a venture fund investing in AI technology companies receives a capital contribution from the dedicated AI account of the fund of funds, that fund is required to invest in the AI sector. When an AI startup attempts to raise investment, approaching such funds increases the probability of opening doors. While existing funds of funds had sector-specific accounts such as biotechnology, cultural content, and materials and components, the AI Basic Act has created the legal basis to add an AI account.

3. AI Research Institutes: Research Hubs Built Jointly by Universities and Companies

Article 22-2 is titled "Establishment and Support, etc. of AI Research Institutes." This provision was newly created by the amended act of January 20, 2026 (Act No. 21311) and took effect on January 22, 2026. If the AI Safety Institute under Article 12 is a government-led safety assessment body, the AI research institute under Article 22-2 is a private sector-led technology research hub. Their characteristics are entirely different.

Let us examine the establishing entities. Entities specified by Presidential Decree, such as universities and companies, may establish them individually or jointly. Article 16-2 of the Enforcement

Decree specifies that scope. Universities under the Higher Education Act, government-funded research institutes in science and technology, non-profit corporations, companies under the Commercial Code, government-funded research institutes in economics, humanities, and social sciences, research institutes funded by local governments, specific research institutes, and corporations or organizations recognized as suitable by the Minister of Science and ICT all qualify. Major conglomerates like Samsung Electronics, universities like KAIST, and private research foundations can all establish AI research institutes. Samsung Electronics and KAIST can also establish one jointly.

There are three authorization requirements: having three or more incorporators, possessing the manpower and facility capabilities to perform statutory projects, and satisfying additional requirements prescribed by the Enforcement Decree. Article 16-2, Paragraph 2 of the Enforcement Decree requires, as additional conditions, financial capability necessary for independent operation, and security measures for protecting research data and preventing personal data leaks. Upon applying for authorization, a three-year research and development business plan and a financial resource procurement plan must be submitted to the Minister of Science and ICT (Article

16-3 of the Enforcement Decree).

Projects conducted by AI research institutes number six: research and development of AI technology, convergence research and development with other technologies and disciplines, management, transfer, utilization, and commercialization of R&D outcomes, training of specialized personnel, international exchanges and cooperation, and other matters necessary for research and development. The third item, "management, transfer, utilization, and commercialization of outcomes," stands out. The law has explicitly designated as a statutory project of the research institute the pathway through which technology originating from the institute is transferred to enterprises and becomes a product entering the market.

Operating financial resources are divided into three streams. The government and local governments may support expenses within the scope of their budgets (Article 22-2, Paragraph 6), subsidies or donations may be received from the private sector (Paragraph 7), and profit-making businesses may be conducted as determined by the articles of incorporation (Paragraph 7). The legal form corresponds to an incorporated foundation under the Civil Act (Paragraph 9).

There are also mechanisms for securing personnel. The head of a research institute may, in consultation with universities, companies, etc., have executives and employees dispatched or hold concurrent positions (Paragraph 8). This means that a KAIST professor can hold a concurrent research position at an AI research institute, or a researcher from a major conglomerate can come to work on dispatch. Necessary support can also be provided to the home institution of dispatched employees.

The Minister of Science and ICT holds supervisory authority. If authorization was obtained through false or fraudulent means, or if achieving the purpose becomes impossible, authorization must be revoked. Corrective orders or revocation of authorization are also possible in cases of failing to meet authorization requirements, conducting business outside designated purposes, violating laws or articles of incorporation, harming public interest, failing to commence operations within six months of authorization, or lacking performance records for over one year.

Article 22-3, newly established by the same amendment, provides the legal basis for the Minister of Science and ICT to directly establish and operate AI research institutions. It carries the condition "when

necessary to secure innovative artificial intelligence technology."

While Article 22-2 is private-led, Article 22-3 is government-led. It is a provision that allows the government to directly establish research institutes in areas where the private sector does not move, such as basic research or AI technology related to national security where market profitability is uncertain.

A comparison with China's National AI Open Innovation Platforms reveals the difference. Since 2017, China has designated private tech giants such as Baidu (autonomous driving), Alibaba (smart cities), Tencent (medical imaging), iFlytek (speech recognition), and SenseTime (computer vision) as national AI platforms, channeling government projects and data to them. It is a structure in which the government selects champion enterprises and nurtures them. South Korea's AI research institute system does not restrict establishing entities. Whether a university, SME, or government-funded research institute, any entity that meets the requirements and obtains authorization can establish a research institute. The direction is decentralized.

4. Securing Specialized Personnel: Domestic Nurturing and Overseas Recruitment

One of the hardest challenges for AI startups is finding talent. Article 21 of the AI Basic Act stipulates a two-pronged talent policy under the title "Securing Specialized Personnel."

Domestic nurturing is under Paragraph 1. Projects that the Minister of Science and ICT may pursue include job capability education and training programs for specialized personnel, employment support and facilitating the influx of new personnel, expanding opportunities to enter public office, domestic and international training and international exchanges, and improving working environments and treatment. The fifth item, "improving working environments and enhancing welfare such as treatment," is a rare expression in statutes. It reflects the reality that the departure of AI personnel is driven not only by salaries, but also by working conditions.

The third item also catches one's attention: "Promoting career diversification, including expanding opportunities for specialized personnel to enter public office." It means broadening pathways so that AI experts do not only go to major conglomerates or startups, but can also enter the government and public institutions. In a situation where vice-minister-level Chief Artificial Intelligence Officers (CAIOs) have been designated across 48 ministries and agencies, practical

working personnel are needed to execute AI-related tasks in each ministry.

Overseas talent recruitment is under Paragraph 2. Starting with investigating and analyzing specialized personnel at overseas universities, research institutions, and enterprises, it lists six policy measures up to building international networks, supporting domestic employment of foreign specialized talent, promoting overseas expansion of domestic research institutions and attracting foreign research institutions domestically, and hosting international organizations and international events domestically.

For this provision to function in practice, policies outside the law must support it. The AI Graduate School Program designated by the government since 2019 forms one pillar. Specialized AI graduate schools were established at KAIST, Korea University, Sungkyunkwan University, POSTECH, UNIST, and others, with more than ten operating as of 2026. The "AI Talent Nurturing Direction," approved at the 2nd Plenary Meeting of the National Artificial Intelligence Strategy Committee on February 25, 2026, constituted one axis of the Republic of Korea AI Action Plan containing 99 implementation tasks.

Let us compare this with the EU. The EU did not include provisions for nurturing specialized personnel within its AI Act. Instead, it supports AI research personnel through the Horizon Europe framework program (2021–2027, totaling 95.5 billion euros) and the Digital Europe Programme (7.5 billion euros). AI regulation and talent development operate on distinct legal bases. South Korea's AI Basic Act encompasses regulation, promotion, and talent development within a single statute. The advantage is that inter-policy coordination is easy, while the disadvantage is that the specificity of individual policies may be weak within the framework of a basic act.

5. Pathways Through Which Public Data Becomes Training Resources

Training AI models requires data. The cost of obtaining good data is a major burden for startups. The AI Basic Act addresses this issue through two pathways.

Article 15 regulates overall policies regarding training data. The Minister of Science and ICT must pursue policies for the production, collection, management, distribution, facilitation of utilization, and quality assurance of training data in consultation with relevant ministries. The government may directly execute training data creation projects (Paragraph 3), and must build an integrated provision and management system (integrated provision system) for

training data to provide it for free use by the private sector (Paragraph 4). Because it states "shall provide," it is a mandatory duty rather than discretionary.

The second pathway is Item 2-2 of Article 16, Paragraph 2, added by the 2026 amended act. It stipulates that public institutions may generate and provide data held and managed by them as training data, and may provide support necessary to ensure appropriate quality levels. This grants a legal basis for processing and providing data held by public institutions for AI training purposes.

The practical significance of this provision is substantial. South Korean public institutions hold vast amounts of data. Examples include health checkup data from the National Health Insurance Service, traffic volume data from the Ministry of Land, Infrastructure and Transport, meteorological observation data from the Korea Meteorological Administration, and patent document data from the Korean Intellectual Property Office. When this data is provided for training purposes, the development costs for medical AI, traffic AI, climate prediction AI, and patent analysis AI drop significantly. Although practical issues such as conflicts with the Personal Information Protection Act, data quality verification, and

de-identification processing remain, the establishing of a legal foundation itself marks the starting point.

Article 12 of the Enforcement Decree specifies selection criteria for projects eligible for training data support. Targets include the development of production and processing technologies for training data, facilitation and quality assurance of data for AI service development, and development of guidelines and standard contracts. Upon selection, economic spillover effects and compliance with laws and regulations are to be considered together.

6. Safety Nets for SMEs and Vulnerable Groups

Article 17 of the AI Basic Act is a short yet powerful provision: "When implementing various support policies related to artificial intelligence technology and the artificial intelligence industry pursuant to this Act, priority consideration shall be given to small and medium-sized enterprises, etc." Small and medium-sized enterprises, venture companies, and micro-enterprises become priority consideration targets across various support measures. Activating SME participation must also be reflected in basic plans (Paragraph 2). The Minister of Science and ICT may support SMEs in implementing measures and impact assessments related to high-impact AI (Paragraph 3). This means the government intends to help offset the

very costs incurred in complying with regulations.

Item 3 of Article 16, Paragraph 2 stipulates educational support related to the adoption and utilization of AI technology for executives and employees of SMEs, venture companies, and micro-enterprises.

Item 4 provides financial support used for their adoption and utilization of AI technology. This serves as the legal basis for the government to directly support funds for AI adoption.

Article 17-2, added by the 2026 amended act, is a provision supporting all or part of the costs for individuals who experience difficulty using AI products and services due to economic conditions. Article 1-2 of the Enforcement Decree defines the scope of vulnerable groups regarding AI into 10 categories. These include recipients of basic living allowances and the near-poor class, senior citizens aged 65 or older, persons with disabilities, women with career breaks, job seekers, farmers and fishermen, members of multicultural families, North Korean defectors, single-parent families, and persons recognized by central administrative agencies or local governments as experiencing difficulty using AI products and services. Usage cost support for vulnerable groups was a mechanism absent from the original law, constituting a safety net created for the first time by the

2026 amendment.

Article 15-2 of the Enforcement Decree further expands the scope of cost support beneficiaries. In addition to nine of the above categories, it includes science and engineering personnel, persons of national merit who have been evaluated with disability grades, and regional balanced talent studying at or graduated from regional universities. A legal pathway has opened for young people from regional universities to receive support for AI service usage costs.

7. Demonstration Infrastructure and Integrated Clusters: The Pathway for Technology to Enter the Market

AI technology originating from research institutes does not immediately enter the market. It must undergo demonstration testing in between. Article 24 allows the state and local governments to construct and operate facilities and equipment to support demonstration, performance testing, and verification/certification of technologies developed or transferred by AI business operators. Demonstration infrastructure held by public institutions may also be opened to business operators.

Article 23 regulates the AI integrated cluster system. It is a structure that designates areas where enterprises, institutions, and

organizations conducting AI research and development are physically concentrated, focusing administrative, financial, and technical support upon them. Dedicated institutions may be established or designated to comprehensively support integrated clusters, and contributions or subsidies may be provided for operating expenses.

The convergence promotion provision in Article 19 is also important to startups. It stipulates that the government must actively support artificial intelligence convergence products and services developed with priority reflection as national R&D project tasks so that temporary authorizations and regulatory sandbox demonstrations under the Special Act on Promotion of Information and Communications Technology and Activation of Convergence, etc. can be smoothly implemented. It provides the legal backdrop for overcoming situations blocked by existing regulations at the intersection of AI with medicine, AI with finance, and AI with transportation through temporary authorization.

8. Expansion of Financial Resources: Where to Source Funds

Article 37 is titled "Expansion of Financial Resources for the Promotion of the Artificial Intelligence Industry, etc." The state must formulate measures to continuously and stably expand financial resources necessary to effectively implement basic plans and policies

pursuant to this Act. The Minister of Science and ICT may recommend public institutions to provide support necessary for AI promotion projects. The state and local governments must prepare necessary measures so that the private sector can actively invest in the AI industry.

This provision is declaratory. It does not specify concrete monetary amounts or fund sizes. The actual scale of financial resources is decided during the annual budget process. The "Republic of Korea AI Action Plan (2026–2028)," finalized at the 2nd Plenary Meeting of the National Artificial Intelligence Strategy Committee on February 25, 2026, contains 99 implementation tasks and 326 policy recommendations; the budget allocated to each task serves as the channel for realizing the duty to expand financial resources under Article 37.

9. What Startups Can Extract From This Law

Summarizing the entirety of industrial promotion provisions from the perspective of startups reveals the bigger picture.

In the phase of searching for investment funds, there is the dedicated AI account of the venture investment fund of funds (Article 18, Paragraph 3). When it is difficult to prove the value of technology,

one can look forward to the government's AI technology valuation system (Article 18, Paragraph 1, Item 4). Education and training during early business stages are supported by startup activation training projects (Article 18, Paragraph 1, Item 2). If one wishes to adopt and utilize R&D outcomes, access is available to AI research and technology development outcomes provided by the government (Article 18, Paragraph 1, Item 5).

When data is lacking, the provision of public data for training purposes (Article 16, Paragraph 2, Item 2-2) and the integrated provision system (Article 15, Paragraph 4) serve as pathways. To nurture or secure research personnel, specialized talent development projects (Article 21) and the personnel dispatch system of AI research institutes (Article 22-2, Paragraph 8) can be utilized. Funds for adopting AI technology (Article 16, Paragraph 2, Item 4) and priority consideration for SMEs (Article 17, Paragraph 1) reduce early cost burdens.

After a product is released, demonstration infrastructure (Article 24), regulatory sandbox demonstrations (Article 19, Paragraph 3), and priority consideration in public procurement (Article 16, Paragraph 3) can assist market entry. Giving priority consideration to AI products

and services in public procurement, alongside an immunity clause (Article 16, Paragraph 4) stating that public officials in charge shall not bear indemnity liability for damages occurring without intent or gross negligence, functions as a mechanism to encourage AI adoption decisions by public officials.

If the cost of fulfilling obligations for high-impact AI poses a burden, there is a provision under which the Minister of Science and ICT supports SMEs in compliance and impact assessments (Article 17, Paragraph 3). It signifies that the government will share regulatory compliance costs.

10. Remaining Questions

Having a legal basis in statute is not synonymous with actual support being delivered. Although Article 18 permits the utilization of venture investment funds of funds, how much will actually be allocated to the dedicated AI account depends on the annual budget and operational plans of Korea Venture Investment Corp. What specific level is required for the substantive authorization requirement of "financial capability necessary for independent operation" in establishing AI research institutes awaits notifications from the Ministry of Science and ICT. How the provision of public data for training will harmonize with the Personal Information Protection

Act and the Copyright Act also remains a matter for follow-up regulatory adjustments.

Under circumstances where the EU has set an investment target of 200 billion euros through InvestAI and China executes AI subsidies at local government levels, while South Korea's legal foundation has been established, the scale of financial resources invested cannot yet be compared.

Still, one thing is clear. To the three CEOs having lunch at that shared office in the autumn of 2025, this law does not merely impose regulatory obligations. Pathways for investment capital, research institutes from which technology can be adapted and utilized, institutional backing for nurturing and recruiting talent, access routes to public data, and priority status for entering government markets—these exist within the law. The question is whether these provisions will actually function. The statutory text has been written, and the Enforcement Decree has been issued. Attaching budgets, refining authorization procedures, and actually opening data—filling the space in between remains a work in progress.

Part 7. When Violations Occur: Sanctions and Risk Management

Chapter 21. Initiation Requirements for Factual Investigations and Responding to Document Submission

Suppose an AI-generated deepfake video is circulated on an online platform. There are dozens of victims, and the video is shared tens of thousands of times across social media. The Ministry of Science and ICT demands document submission from the AI service provider of the platform. Although it is during the grace period, an investigation is initiated. This is the moment when Article 40 of the AI Basic Act comes into play.

Factual investigation is the sharp edge of this law. Although there is criticism that the upper limit of the administrative fine—30 million KRW—is virtually nominal compared to the EU's maximum limit of 35 million EUR, it is not the amount of the administrative fine that inflicts a substantial blow on a business operator. The moment an investigator steps into the office and says they will inspect books and documents, the world of the business operator changes.

1. Legal Structure of Factual Investigations

Article 40 of the AI Basic Act contains three paragraphs under the title "Factual Investigations, etc."

Paragraph 1 establishes the initiation requirements. There are two cases in which the Minister of Science and ICT may initiate a factual investigation. One is when a violation is discovered or when suspicion of a violation comes to light. The other is when a report or a civil complaint regarding a violation is received.

There are five provisions subject to investigation for violations: the obligation to label generative AI outputs (Article 31, Paragraph 2); the obligation to provide clear notification or labeling for outputs such as deepfakes that are difficult to distinguish from reality (Article 31, Paragraph 3); the obligation to identify, evaluate, and mitigate risks of high-performance AI (Article 32, Paragraph 1); the obligation to submit the results of such implementation (Article 32, Paragraph 2); and the obligation of high-impact AI business operators to take measures to ensure safety and reliability (Article 34, Paragraph 1).

What is notable is what has been omitted. A violation of the prior notification obligation under Article 31, Paragraph 1 is not subject to factual investigation. A violation of the prior notification obligation is subject to the direct imposition of administrative fines under Article 43. This means an administrative fine can be imposed even without a request for document submission or an on-site inspection.

Conversely, violations of obligations to label outputs or to ensure the safety of high-performance AI are structured such that administrative fines are not imposed immediately; instead, if a violation is confirmed following a factual investigation, a suspension order or corrective order is issued. Only when that order is not complied with will an administrative fine be attached.

Paragraph 2 defines the means of investigation. Affiliated public officials may enter the offices or business places of a business operator to inspect books, documents, and other materials or items. Matters not specified in this Act shall be governed by the Framework Act on Administrative Investigations.

Paragraph 3 specifies the follow-up measures. If a violation is recognized, necessary measures may be ordered for the suspension or correction of the relevant violation.

Compressing this structure into a single sentence yields the following: the recognition of suspicion or the receipt of a report serves as the triggering requirement; demands for document submission and on-site entry inspections serve as the investigative means; suspension orders or corrective orders serve as sanctions; and failure to comply with such orders serves as the prerequisite for

an administrative fine of up to 30 million KRW.

2. Grace Period and Exceptions: When Is the Sword Drawn?

Although the AI Basic Act came into force on January 22, 2026, there is a grace period for the imposition of administrative fines and factual investigations. This is not a deferral specified in the Act or the Enforcement Decree, but an administrative enforcement policy announced by the Ministry of Science and ICT. The period is at least one year or more, and in the early stages of implementation, administrative guidance and corrective recommendations are prioritized. The intention is to choose guidance over punishment for violations resulting from business operators being unaware of obligations or not knowing how to fulfill them.

The grace period does not apply to all cases. The government stated that in extremely exceptional cases where severe social issues arise—such as casualties or human rights violations—it will conduct factual investigations even during the grace period.

What situations fall into this category? A representative case is when circumstances indicate that AI generation technology has been used in crimes such as the production of sexual exploitation material, defamation, or fraud. The same applies when a malfunction or

misjudgment of an AI system directly inflicts harm on human life or body. Cases where an accident occurs due to a misjudgment of an autonomous driving assistance system, or where a misdiagnosis by an AI medical device causes physical harm to a patient, belong here. It is also the case when an AI system leads to systematic discrimination in high-impact domains such as hiring or loan screening.

The common denominator in these exceptional situations is the severity and urgency of the damage. While business operators can be given time to correct omissions of transparency obligations or inadequate documentation, there is no reason to wait when human lives are at stake or crimes are ongoing.

From the perspective of a business operator, this is a clear boundary line. While there is a buffer zone called the grace period, that buffer zone is perforated with holes.

3. When Investigators Arrive: Procedures of Factual Investigations

Once a factual investigation is initiated, the Minister of Science and ICT uses two pathways.

One is a request for document submission. An official document demanding the submission of relevant materials is sent to the business operator. The other is an on-site inspection. Affiliated public officials enter the office or business place of the business operator.

Article 40, Paragraph 2 of the AI Basic Act states that except for matters specified in this Act, the contents, methods, and procedures of investigations shall be governed by the Framework Act on Administrative Investigations. This delegation carries two implications for business operators.

The Framework Act on Administrative Investigations is equipped with procedural safeguards to protect business operators. According to Article 17 of the Framework Act on Administrative Investigations, administrative agencies must provide prior written notice to the subject of investigation at least seven days before the initiation of the investigation. Pursuant to Articles 9 through 11 of the same Act, summons for appearance, requests for document submission, and notices of on-site entry inspection must state the purpose, period, and scope of the investigation, the names and positions of the investigators, the materials to be submitted, and sanctions for refusing investigation along with the underlying statutes. In principle,

on-site inspections cannot be conducted before sunrise or after sunset, with exceptions existing when the subject of investigation consents or when conducted during business hours of the workplace. Investigators must carry identification indicating their authority and present it to the subject of investigation.

However, there are exceptions to prior notification. If it is deemed that prior notification poses a risk of evidence destruction, an inspection notice may be presented at the time of initiating the investigation, or the purpose of the investigation may be notified orally. There is room for this exception to apply in matters where evidence can disappear rapidly, such as deepfake crimes.

The basic principles set forth in Article 4 of the Framework Act on Administrative Investigations also apply. Investigations must be conducted within the minimum extent necessary to achieve their purpose. Investigative powers must not be abused for other purposes. Pursuant to Article 15 of the same Act, the same subject of investigation cannot be re-investigated for the same matter, except when new evidence suspecting illegal acts has been secured. Information learned through the investigation must not be used for purposes other than the original investigative purpose or provided to

third parties.

The final provision is important for business operators. The law blocks information acquired by the Ministry of Science and ICT during the investigation process from leaking to other agencies, being exposed to competitors, or being disclosed to the press. Article 4, Paragraph 5 of the Framework Act on Administrative Investigations stipulates that the subject or content of an administrative investigation shall not be made public, nor shall secrets learned in the course of duty be leaked.

4. Scope and Limits of Document Submission

In a factual investigation, what business operators react to sensitively is the request for document submission. The statutory phrasing "may require submission of relevant materials" is broad. How far does the scope of materials that a business operator must present extend?

The scope is determined by the provisions subject to investigation under Article 40. If the suspicion involves a violation of the obligation to label outputs, the labeling mechanism of the relevant generative AI service, watermark application history, screenshots of notification screens provided to users, and deepfake filtering logs become subject

to document submission. If the suspicion involves a violation of the obligation to ensure the safety of high-performance AI, risk identification reports, risk evaluation results, mitigation measure implementation records, and risk management system documents are subject. If the suspicion involves a violation of obligations of high-impact AI business operators, risk management plans, explanation plans, user protection plans, human oversight and supervision systems, and implementation evidence documents (obligatory 5-year retention) are subject.

Even if the scope appears broad, there are limits.

The principle of proportionality established by Article 4, Paragraph 1 of the Framework Act on Administrative Investigations serves as the first limit: the minimum extent necessary to achieve the investigative purpose. Comprehensively demanding overall business materials unrelated to the suspicion violates the principle of proportionality. Demanding a company's total revenue data or personnel records in an investigation into a violation of deepfake labeling obligations exceeds the scope of the principle of proportionality.

The items required to be stated in the request for document submission constitute the second limit. According to Article 10,

Paragraph 2 of the Framework Act on Administrative Investigations, a request for document submission must state the submission period, reason for request, requested documents, whether submitted documents will be returned, sanctions for refusal of submission, and other necessary matters related to the relevant administrative investigation. There is no obligation to submit materials that are not stated in the request. Upon receiving the request, business operators must thoroughly verify the stated items. If the scope of submission is deemed excessive, there is grounds to request a reduction to the minimum extent necessary in light of the purpose of the investigation.

5. Protection of Trade Secrets and Technical Secrets

The core assets of an AI business operator are model architecture, training dataset composition, hyperparameters, and prompt engineering know-how. If this information is included in the scope of submission during a factual investigation, the business operator's competitiveness can collapse in an instant.

There is no provision directly stipulating the protection of trade secrets in the statutory text of the AI Basic Act. However, the proviso to Article 27, Paragraph 1 of the Enforcement Decree specifies that high-impact AI business operators may exclude matters

corresponding to trade secrets when posting risk management plans and other details. In the context of a factual investigation, grounds for protection derive from the Framework Act on Administrative Investigations and other statutes.

Article 4, Paragraphs 5 and 6 of the Framework Act on Administrative Investigations form the first line of defense. Investigation contents must not be made public, secrets learned in the course of duty must not be leaked, and information learned through investigation must not be used for purposes other than the original investigative purpose or provided to third parties. This constitutes an individual duty of confidentiality for investigators as well as an agency-level duty to restrict the purpose of information.

The Unfair Competition Prevention and Trade Secret Protection Act serves as the second line of defense. Article 2, Item 2 of this Act protects as trade secrets technical or managerial information useful for business activities—such as production methods and sales methods—that is not publicly known, possesses independent economic value, and has been maintained as secret. AI model weights, training data curation methods, and reinforcement learning reward function designs can fall under this definition. If a business

operator has managed this information as secret and it possesses economic value, it receives legal protection as a trade secret.

In practice, there are three responses a business operator can take.

The first is to submit materials while explicitly marking information that constitutes trade secrets and requesting confidentiality. If materials falling under trade secrets exist within the scope stated in the request, the business operator may substantiate that the materials are trade secrets and request to restrict the scope of inspection. There is no provision in this Act directly penalizing the refusal of document submission itself. Nonetheless, uncooperative behavior during an investigation works disadvantageously in determining whether a violation occurred; furthermore, if a suspension order or corrective order under Article 40, Paragraph 3 is issued based on investigation results, non-compliance leads to administrative fines under Article 43, Paragraph 1, Item 3. Thus, negotiating scope adjustments is more realistic than outright refusal.

The AI Basic Act Support Desk plays a role at this point. Operated by the Korea Software Industry Association (KOSA), this desk provides anonymous consulting. A business operator facing a request for document submission can receive consultations on where to draw

the line between trade secrets and the obligation to cooperate with investigations. In-depth legal reviews take up to 14 days.

6. EU Market Surveillance Authorities: A Wider Sword, Heavier Penalties

The European Union AI Act grants powers corresponding to factual investigations to market surveillance authorities. This structure differs from South Korea, where the Ministry of Science and ICT exercises investigative authority independently.

According to Article 70 of the EU AI Act, each member state must designate one or more market surveillance authorities. These authorities are responsible for market surveillance of high-risk AI systems. The AI Office holds exclusive supervisory authority over general-purpose AI (GPAI) models. Supervision by the AI Office over business operators developing both GPAI models and AI systems based on them was present in the original regulation, and the Digital Omnibus amendment—proposed in November 2025 and entering into force in July 2026—further clarified that jurisdiction. AI systems integrated into very large online platforms (VLOPs/VLOSEs) also came under the supervision of the AI Office through this amendment.

The investigative powers of market surveillance authorities are broader than those in South Korea. Pursuant to Article 74, they can

access technical documentation, conformity assessment records, and training data-related records; if verification is difficult through other means, they may also access the source code of AI systems. Pursuant to Article 79 et seq., they can conduct testing of AI systems if necessary, and order the withdrawal or recall of AI systems from the market.

The scale of sanctions is decisively different. Fines under the EU AI Act vary according to the type of violation. Violations of prohibited AI practices carry a maximum limit of up to 35 million EUR or 7% of total worldwide annual turnover, while violations of other obligations carry a limit of 15 million EUR or 3%. This is incomparable to South Korea's 30 million KRW.

South Korea looks at impacts. The European Union looks at risks. This distinction permeates directly through to the enforcement stage. Factual investigations in South Korea focus on verifying violations and ordering corrections. Market surveillance in the EU goes as far as removing dangerous AI systems from the market. In South Korea, the end of a factual investigation is a corrective order, and disobeying that order results in a 30 million KRW administrative fine. In the EU, the end of market surveillance is a market withdrawal order, and

violating it leads to tens of billions of KRW in fines.

China accelerates in a different manner. The Cyberspace Administration of China (CAC) imposes a prior filing obligation (备案) on generative AI services, securing access to technical materials during the filing review process. Internet information offices across various regions have cleared generative AI apps and mini-programs that have not undergone security assessments; in June 2025, Shanghai authorities deleted 820,000 pieces of illegal content and deactivated approximately 2,700 AI agents. Rather than post-hoc investigation, prior filing and constant monitoring serve as the axis of enforcement. While violations of the Personal Information Protection Law can carry fines of up to 50 million RMB (approximately 10 billion KRW) or up to 5% of the preceding year's turnover, practically stronger sanctions are service suspension and business shutdown.

Placing the investigative powers of the three jurisdictions on a single line yields the following: South Korea investigates upon discovering suspicions and orders correction; the EU operates a market surveillance system continuously and removes non-compliant systems from the market; China catches issues in advance through prior filing and constant monitoring, and terminates

services after the fact.

7. Business Operators' Preparation: Before the Investigation Arrives

To prepare for the possibility of undergoing a factual investigation, business operators must check four items.

First is the organization of high-impact AI determination records. Business operators must maintain records of prior reviews assessing whether their AI falls under any of the eleven domains listed in Article 2, Item 4. Even if an AI is determined not to be high-impact, it is safer to document the grounds for that determination. Article 33 imposes an obligation on business operators to review in advance whether their AI constitutes high-impact AI.

Second is the retention of implementation evidence documents. Article 34, Paragraph 1, Item 5 requires the preparation and retention of documents that can verify the details of measures taken to ensure safety and reliability. The Enforcement Decree sets the retention period at five years. It is highly likely that these documents will be requested first in a factual investigation.

Third is the inspection of output labeling systems. Since factual investigations can be initiated even during the grace period for

matters that have become major social issues like deepfake crimes, business operators providing generative AI services must verify their compliance status with labeling obligations—whether visible labels exist, whether watermark application logs remain, and whether the manner of notifying users of AI generation satisfies legal requirements.

Fourth is the establishment of a trade secret management system. Management measures as trade secrets must be in place for model architecture, training data composition, and core algorithms. Access controls, confidentiality markings, and non-disclosure agreements must be established to claim protection for such materials during a factual investigation. Information not managed as secret will not receive trade secret protection.

There is one reason why this preparation is necessary: the grace period will end. Afterward, factual investigations will be possible for violations of obligations enumerated in Article 40, and administrative fines will be imposed if corrective orders are disobeyed. The gap between business operators who establish systems during the grace period and those who rush to move only after it ends will be revealed on the day the first notice of factual investigation arrives.

A factual investigation is not a punishment. Article 4, Paragraph 4 of the Framework Act on Administrative Investigations states: "Administrative investigations shall focus on encouraging compliance with statutes and regulations rather than punishing violations of statutes and regulations." Is a factual investigation a tool of detection and punishment, or a process of encouraging compliance? South Korean law has chosen the latter. That choice is consistently reflected in the administrative fine upper limit of 30 million KRW, the grace period of at least one year, and the administrative fine structure premised on non-compliance with corrective orders. The message this law sends to business operators is not to fear fines, but to build a system.

행정법 제4조 제4항, 제5조 제1항. 제5조 제2항
제5조 제3항. 제5조 제4항

Chapter 22. Administrative Fines, Corrective Orders, Penalties Framework, and the Reality of Grace Periods

In March 2026, a phone call from a startup CEO came into the AI Basic Act Support Desk. The CEO stated that their company had been operating a chatbot service for two months without inserting the prior notification statement for AI. "How much is the administrative fine?" The counselor read the legal provision: 5 million KRW for a first violation. After a brief silence, the CEO asked, "But we are currently in a grace period, so it won't be imposed, right?" The counselor's answer was brief: the legal obligation had already taken effect as of January 22, and a grace period does not imply legal immunity.

This single phone call encapsulates the essence of the AI Basic Act's enforcement framework. The law is in effect, obligations are active, and administrative fine provisions exist. However, the government publicly declared that it will prioritize guidance over punishment for at least one year. How a business operator interprets this gap will determine their response.

1. Three Administrative Fines, Three Degrees of Weight

Article 43(1) of the AI Basic Act limits the targets of administrative fines to three cases. Considering the entire law consists of a six-chapter structure ending with Article 43, this means there are only

three violations for which administrative fines can be imposed.

Failing to comply with the prior notification obligation is the first case. Article 31(1) requires that when providing products or services utilizing high-impact AI or generative AI, users must be informed in advance that "this service is operated based on artificial intelligence." Omitting this results in an administrative fine under Article 43(1)(1). According to the individual criteria in Table 2 of the Enforcement Decree, the fines are 5 million KRW for the first violation, 10 million KRW for the second violation, and 15 million KRW for the third or subsequent violations. This is the lowest amount among the three administrative fines. Prior notification is not technically difficult. It simply requires adding a single line of text on a service screen. Thus, the fine appears to have been set low accordingly.

Failing to designate a domestic agent is the second case. Pursuant to Article 36(1), overseas AI business operators without an address or place of business in Korea whose number of users or revenue exceeds the threshold set by the Enforcement Decree must designate a domestic agent in writing and report to the Minister of Science and ICT. In the event of non-designation, the administrative fine is fixed at 20 million KRW for all instances from the first to the third violation.

There is no differentiation based on the number of violations. This signals that legislators view evasion of domestic jurisdiction by overseas business operators as a serious matter. This is because a situation where a global AI company secures millions of users in the Korean market without maintaining a domestic point of contact makes law enforcement itself impossible.

Failing to comply with a corrective order is the third case, and carries the heaviest amount. Pursuant to Article 40(3), if a business operator fails to comply with a suspension order or corrective order issued by the Minister of Science and ICT, fines of 10 million KRW for the first violation, 20 million KRW for the second violation, and 30 million KRW for the third or subsequent violations are imposed. At the third or higher violation, it reaches the statutory maximum of 30 million KRW. Non-compliance with a corrective order represents disregard for an official order of an administrative agency, and is thus structured to receive the highest penalty.

2. Mitigation and Aggravation: The Actual Imposed Amounts Vary

The general criteria in Table 2 of the Enforcement Decree include provisions for reducing or increasing the amount of administrative fines.

There are four grounds for mitigation: where the business operator is a small or medium enterprise (SME), venture business, or micro-enterprise under the Small and Medium Enterprises Basic Act; where the violation was caused by minor negligence or error; where efforts to correct or resolve the state of violation are recognized; or where reduction is deemed necessary considering the type and severity of the violation. The scope of mitigation is within one-half of the amount specified in the individual criteria. For example, the base administrative fine of 5 million KRW for a first violation of the prior notification obligation can be reduced down to 2.5 million KRW if the operator is an SME. However, mitigation is not applied to business operators who are in arrears with administrative fines.

There are three grounds for aggravation: where the nature and severity of the violation are severe and recognized to cause significant harm to users; where the duration of the violation state is six months or longer; or where an increase is deemed necessary considering the severity, motive, and consequences of the violation. The scope of aggravation is also within one-half, but cannot exceed the statutory maximum of 30 million KRW. If aggravation is applied to a first violation of failing to comply with a corrective order (base amount of 10 million KRW), it can rise up to 15 million KRW, while a

third violation (base amount of 30 million KRW) is already at the maximum limit, so it remains 30 million KRW even with aggravation.

There are also rules for calculating the number of violations. An aggravated tier applies only when an administrative fine disposition was received for the same violation within the preceding two years. After two years, the violation count resets.

3. Fact-Finding Investigations and Corrective Orders: Administrative Measures Prior to Administrative Fines

Administrative fines are the last resort for violations of output labeling, securing safety, and high-impact AI business operator duties. Positioned ahead of them are two stages: fact-finding investigations and corrective orders.

Article 40(1) prescribes two cases where the Minister of Science and ICT may conduct a fact-finding investigation. One is when a violation or suspected violation is discovered regarding output labeling obligations (Article 31(2) and (3)), safety assurance obligations (Article 32(1) and (2)), or high-impact AI business operator duties (Article 34(1)). The other is when a report or complaint is received regarding violations of these provisions.

The authority of a fact-finding investigation includes demanding the submission of materials and conducting on-site investigations.

Pursuant to Article 40(2), affiliated public officials may enter the offices or business premises of business operators to inspect books, documents, and other materials or items. The details, methods, and procedures of the investigation follow the Administrative Investigations Basic Act.

If a violation is confirmed as a result of the investigation, necessary measures may be ordered to suspend or correct the violation pursuant to Article 40(3). This is a corrective order. The legal effect of a corrective order is an administrative disposition. Only if a business operator fails to comply with this order does it become subject to administrative fines.

There is one point worthy of note. Violations of the prior notification obligation (Article 31(1)) and failure to designate a domestic agent (Article 36(1)) are not subject to fact-finding investigations but are directly subject to administrative fines. Conversely, output labeling obligations (Article 31(2) and (3)), safety assurance obligations (Article 32), and high-impact AI business operator duties (Article 34) follow a three-stage procedure: fact-finding investigation, corrective

order, and then administrative fines upon non-compliance. The reason these two obligations have a direct fine path is interpreted to be because whether a violation occurred is easy to confirm objectively.

A corrective order can be appealed. Because it is an administrative disposition, it can be contested through administrative appeals or administrative litigation. There is a separate appeal process for administrative fine imposition. Under the Act on the Regulation of Violations of Public Order, if an objection is raised within 60 days, trial proceedings take place in a competent court.

4. Penalties: The Sole Act Subject to Criminal Punishment

Article 42 is the sole criminal penalty provision in the entire AI Basic Act.

If a member of the National AI Strategy Council leaks secrets learned in the course of duty to others or uses them for purposes outside their duty, they shall be punished by imprisonment for up to three years or a fine not exceeding 30 million KRW. The legal basis is the confidentiality obligation under Article 7(9).

The Strategy Council is the core deliberation and resolution body for national AI policy. Chaired by the President, it has a total membership limit of 60, comprising 18 ex-officio members and up to 41 civilian members. They handle sensitive information such as national AI strategies, business operator regulatory directions, and criteria for determining high-impact AI. If a member leaks this information, they face criminal punishment regardless of whether they are a public official.

Article 41 contains a provision deeming non-public officials as public officials for the application of penalties. Among council members, non-public official members are deemed public officials when applying Articles 129 through 132 (Bribery) of the Criminal Act. Executives and employees of institutions or organizations delegated with tasks are additionally subject to Article 127 (Divulgence of Official Secrets) of the Criminal Act. This structure ensures that civilian members incur the same criminal liability as public officials if they accept or offer bribes, and if they leak official secrets, they are directly punished under Article 42 of this Act rather than the Criminal Act.

The fact that the only criminal punishment directly established as a penalty under the law is secret divulgence illustrates the character of the AI Basic Act. Administrative fines (administrative penalties) are imposed for violations of business operator duties, whereas criminal penalties are used solely to protect the internal secrets of governance bodies. Even if a business operator omits prior notification, refuses a corrective order, or fails to designate a domestic agent, they will not go to prison. It ends once the administrative fine is paid.

5. Grace Period: Legal Obligations Remain Alive

The expression "grace period" does not appear in the text of the statute. The word "grace period" is not present in the Enforcement Decree either. It is a policy measure announced by the Ministry of Science and ICT simultaneously with the law's entry into force on January 22, 2026. Fines will be deferred for at least one year, until after January 2027, during which administrative guidance and recommendations for correction will take precedence.

The legal basis for this policy is the discretionary authority of the Minister of Science and ICT. Because whether and when to impose administrative fines falls under the discretion of administrative agencies, the government's declaration that it "will not impose administrative fines for the time being" is not unlawful in itself.

However, it is dangerous for business operators to interpret this as "we do not need to comply with obligations because it is a grace period." Legal obligations have already taken effect as of January 22, 2026. Days without prior notification are accumulating day by day. Even after the grace period ends, the Ministry of Science and ICT can impose administrative fines based on violations that occurred during that period. This is because under the Act on the Regulation of Violations of Public Order, administrative fines can be imposed within five years from the date the violation ceased.

There are exceptions even during the grace period. Immediate fact-finding investigations are possible in serious matters such as deepfake crime exploitation, loss of life, or large-scale human rights violations. The grace period is a deferral for minor procedural deficiencies, not an absolution for severe harm.

In practice, the grace period sends a dual message to business operators. A message to build a compliance framework because the law has been enacted, and a message that there is no need to rush because penalties are not yet being enforced, are heard simultaneously. This ambiguity is one of the reasons why 552 consultation requests flooded the AI Basic Act Support Desk within

just ten weeks of implementation. Transparency-related topics accounted for the largest share at 51%, followed by high-impact AI classification at 19.6%. Business operators were experiencing confusion over "what must be done right now" rather than the obligations themselves.

6. What Changes When the Grace Period Ends

The exact end date of the grace period has not yet been officially announced. Since the Ministry of Science and ICT stated "at least one year," it is unlikely to end before January 22, 2027. However, it will not automatically end after one year either. The government is expected to announce the end date separately.

Once the grace period ends, the imposition of administrative fines for failing to fulfill prior notification obligations will become a reality. The initiation of fact-finding investigations is also highly likely to become active. Administrative activities of issuing corrective orders and tracking compliance are expected to begin in earnest.

From a business operator's perspective, the scenario requiring the utmost caution is this: immediately after the grace period ends, the Ministry of Science and ICT launches a comprehensive inspection. Business operators without prior notification notices on their service

screens, those failing to label AI-generated content, and those lacking high-impact AI risk management documentation could be caught all at once. At that point, defenses such as "we did not know" or "we were preparing" may be recognized as mitigating factors, but they cannot erase the violation itself.

7. Is 30 Million KRW Enough?

The maximum administrative fine under the AI Basic Act is 30 million KRW. That is 30 million South Korean Won, which is approximately 20,000 USD.

The maximum administrative fine under the EU AI Act is 35 million EUR or 7% of total worldwide annual turnover, whichever is higher. Converted at the exchange rate as of July 2026, this amounts to approximately 59.5 billion KRW. Between Korea's 30 million KRW and the EU's 35 million EUR, there is a gap of more than 1,000 times.

China applies fines under relevant laws such as the Cybersecurity Law and Data Security Law for violations of sector-specific AI regulations, while concurrently ordering business suspensions or service shutdowns. The revised Cybersecurity Law, which took effect on January 1, 2026, significantly raised the fine caps. Observations suggest that the fine upper limits will be even higher in the

comprehensive AI law China is currently preparing. Korea's 30 million KRW is lower even than China's current fine levels.

Placing the three jurisdictions side-by-side renders the ranking of enforcement severity sharp and clear. The EU's 35 million EUR stands at the top, China—which adds business suspension on top of fines—occupies the middle, and Korea's 30 million KRW sits at the bottom. Once China passes its comprehensive law, the gap is likely to widen further.

Questions regarding whether an amount of 30 million KRW can provide compliance motivation to large enterprises have existed since the enactment of the law. When administrative fine levels became a bone of contention during deliberations in the Science, ICT, Future Planning, Broadcasting and Communications Committee of the 22nd National Assembly, the rationale that excessive penalties would stifle innovation during the early growth stage of the domestic AI industry justified the low amount. For a company with annual revenues in the trillions of KRW, 30 million KRW is close to a parking ticket. If the profit gained from a violation exceeds the fine, continuing the violation while paying the fine becomes a rational choice.

There is a counterargument. Korea does not rely solely on administrative fines for regulation. Fact-finding investigations, corrective orders, administrative guidance, and guideline recommendations are positioned ahead of administrative fines. Business operators that reject recommendations from the Ministry of Science and ICT and disregard even corrective orders will bear indirect costs such as reputation loss within the industry, exclusion from public procurement, and tightened sanctions through subsequent legislation. This argument posits that the 30 million KRW figure is not everything.

However, indirect costs do not constitute legal enforceability. The EU alters corporate behavior through the sheer figure of 7% of revenue. China compels compliance through threats to business continuity via operational suspensions. Korea does not yet possess such tools.

8. EU: Three-Tiered Administrative Fines Differentiated by Violation Type

The fine system under the EU AI Act is divided into three tiers according to the severity of the violation.

Violations of prohibited AI practices (Article 5) carry the heaviest penalty: up to 35 million EUR or 7% of total worldwide annual

turnover, whichever is higher. Practices such as social scoring, subliminal manipulation AI, and untargeted biometric surveillance fall under this category. Korea lacks such baseline prohibition provisions, meaning there are no violation acts in place to which fines of this level would even apply.

Violations of high-risk AI obligations and GPAI obligations carry penalties of up to 15 million EUR or 3% of total worldwide annual turnover.

Providing inaccurate information to authorities carries penalties of up to 7.5 million EUR or 1% of total worldwide annual turnover. For SMEs, the lower of the respective amounts applies.

Compared to Korea's single cap of 30 million KRW, the EU has differentiated amounts based on violation types. This reflects the principle of proportionality, pairing high penalties with high-risk violations and low penalties with procedural violations.

Through the Digital Omnibus amendments proposed by the Commission in November 2025 and finalized in the first half of 2026, the EU delayed the application date of standalone high-risk AI obligations to December 2, 2027. High-risk AI embedded in products

will apply from August 2, 2028. Only the deadline was pushed back; penalty levels remained unchanged. Although serving a similar function to Korea's grace period, a key difference is that the EU altered the enforcement date itself through legislative amendment, whereas Korea suspended enforcement through administrative discretion while leaving the statute intact.

9. China: Business Suspension More Terrifying Than Fines

AI enforcement in China does not consist of monetary fines alone.

According to the Interim Measures for the Management of Generative Artificial Intelligence Services (effective August 2023), service providers that fail to undergo prior security evaluations or algorithm filings (备案) receive warnings, public criticisms, orders for correction within a specified period, and temporary suspension of relevant services. In cases of concurrent violations of the Data Security Law or Cybersecurity Law, fines under those laws may be accompanied by business suspension orders, revocation of business licenses, and service shutdowns; the revised Cybersecurity Law effective January 1, 2026, significantly increased fine ceilings. In 2025, under the 'Qinglang' cleanup initiative, the Cyberspace Administration of China (CAC) successively removed generative AI apps and mini-programs that had not undergone security

evaluations. In June 2025, during Shanghai's 'Liangjian Fujiang 2025' campaign, 54 non-compliant generative AI apps were deleted and 3 cases were criminally filed.

Korea's AI Basic Act contains no business suspension or service shutdown orders. Corrective orders can be issued, but because the phrasing is "measures necessary for the suspension or correction of the relevant violation" (Article 40(3)), it differs from suspending the business itself.

Regarding the comprehensive AI law China is preparing, an official draft has not yet been publicly released. Speculation suggests it will follow a direction of raising fine caps while mandating training data source logging and safety impact assessments for large models.

Comparing enforcement mechanisms across the three jurisdictions reveals structural differences. The EU blocks economic incentives for non-compliance through massive fines. China adds threats to business survival via operational suspensions and service shutdowns alongside fines. Korea relies entirely on administrative fines and corrective orders, with low maximum caps.

10. Regulatory Review and Potential for Follow-Up Legislation

Article 33 of the Enforcement Decree imposes a regulatory review duty on the Minister of Science and ICT. As of January 1, 2026, every three years, the Minister must review the validity of criteria for high-performance AI safety assurance obligations and business operators subject to domestic agent designation, taking corrective measures accordingly. However, administrative fines are not included in this review scope. Penalty standards can only be changed by amending the Enforcement Decree's attached table, while the 30 million KRW cap can only be changed by amending Article 43 of the Act.

In the March 2025 issue of Issues and Perspectives (No. 2327), Legislative Research Officer Jun-hwa Jung of the National Assembly Research Service pointed out the need to differentiate obligations by AI business operator type. Senior Research Fellow Won-jun Jung of the Korea Legislation Research Institute also reportedly expressed the view in a September 2025 report that enforcement levels are insufficient to guarantee regulatory effectiveness. The 30 million KRW cap was the product of political compromise at the time of legislative enactment, and voices from both academia and policy research institutes note that upward adjustment will be inevitable as the industry matures.

Legislative trends in the EU and China also influence Korea. Although the EU postponed the enforcement of high-risk AI obligations to December 2027 via Omnibus amendments, fine levels were maintained. If China passes its comprehensive AI law, Korea is likely to remain as the sole jurisdiction with noticeably low penalty levels. Pressure for international regulatory harmonization is likely to accelerate the upward revision of Korea's administrative fines.

11. What Business Operators Must Do Right Now

A grace period is a deferral, not an exemption. Legal obligations are already in force, and violation facts remain on record.

Action must begin with prior notification. For services utilizing high-impact AI or generative AI, business operators must notify users before they use the service so that users are aware the service operates based on AI. Merely placing a notice on the screen can prevent a 5 million KRW administrative fine for a first violation.

Output labeling must be audited. AI creation labels must be placed on outputs from generative AI, and for outputs at the deepfake level that are difficult to distinguish from reality, notifications or labels must be displayed in a manner that users can clearly recognize. Invisible watermarks alone are insufficient.

Overseas business operators must verify whether they have designated a domestic agent. Failure to designate results in an administrative fine of 20 million KRW regardless of the number of violations.

Business operators operating high-impact AI must establish risk management documentation, explanation plans, user protection measures, and management oversight systems, and retain relevant documents for five years. Once fact-finding investigations begin after the grace period ends, the presence or absence of these documents will serve as key evidence determining whether a corrective order is issued.

30 million KRW may be a small sum for large corporations. However, a history of administrative fines, records of corrective orders, and administrative investigation histories impact participation in public procurement, expansion into overseas markets, and attraction of investment. Records weigh heavier than numbers.

The grace period will end. Between business operators that complete preparations before it ends and those that move only after it ends, a single administrative fine will make the difference. The administrative record left by that single fine will endure longer than

expected.

□□□ □□□□ □□□ □□□□ □□, □□□ □□ □□. □□ □□□
□□□ □□ □□ • □□ • □□□

Chapter 23. Document Management in Preparation for Supervision — What and How Much Should Be Retained?

In April 2026, a complaint regarding a loan screening AI was filed at a fintech company in Gangnam, Seoul. An office worker in their 30s with a credit score of 750 was denied a loan, while another applicant with the exact same score was approved. The Ministry of Science and ICT (MSIT) launched a fact-finding investigation pursuant to Article 40 of the AI Basic Act. The first document the investigator requested from the company was a record documenting the basis for the AI model's determination. The company produced a risk management plan document. It was a two-page file hastily put together to coincide with the law's enforcement in January 2026. It contained no details on risk identification, no overview of training data, and no records of objection processing. The investigator asked, "Where are the implementation evidence documents required under Article 27(2) of the Enforcement Decree?" The company's response was, "We are still organizing them."

1. The Structure of Documents Required by Law

Article 34(1) of the AI Basic Act requires high-impact AI business operators to implement six measures. The fifth subparagraph pertains to documentation. The text of the provision reads as follows:

"The preparation and retention of documents capable of verifying the contents of measures for securing safety and reliability." The term "document" referred to in the law means a record proving the entire scope of safety and reliability measures. It is not limited to recording risk management alone. It means creating and retaining documents encompassing all of the remaining five items under Article 34(1)—namely, risk management plans (Subparagraph 1), explanation plans (Subparagraph 2), user protection plans (Subparagraph 3), human management and supervision systems (Subparagraph 4), and matters deliberated and resolved by the Committee (Subparagraph 6).

The retention period is specified in Article 27(2) of the Enforcement Decree. It is five years. It is noted in parentheses that retention through electronic methods is included. This is not a requirement to print documents on paper and store them in a cabinet. Electronic files, databases, and cloud storage are all permitted.

The starting point for calculating the five-year period is not separately specified in the Enforcement Decree. The safest interpretation is to count from the time the measure is implemented. If a risk management plan was established in March 2026, it should be

retained until March 2031. If the plan is amended in the interim, it is safe to retain both the pre-amendment version and the post-amendment version for five years each. This structure makes version control effectively mandatory.

2. Six Types of Documents Subject to Retention

Let us examine the documents corresponding to each subparagraph of Article 34(1) one by one.

The risk management plan document (Subparagraph 1) covers the entire process of risk management policy, organizational structure, and risk identification, analysis, evaluation, treatment, and improvement. Article 27(1)(1) of the Enforcement Decree mandates posting "major contents of the risk management plan, such as risk management policies and organizational structures," on the homepage. Separately from the summary for posting, internal retention documents must include a list of identified risks, evaluation scores by risk, implementation schedules for mitigation measures, and the rationale for judging residual risks.

The explanation plan document (Subparagraph 2) records, to the extent technically feasible, the final results derived by the AI, the primary criteria used to derive those results, and an overview of the

training data used for development and training. The inclusion of the proviso "to the extent technically feasible" does not mean that every single inference path of a black-box model must be fully detailed. However, it is safe to record items such as results from applying explainability tools like SHAP and LIME, weight information of major input variables, and the source, scale, and bias test results of training data.

The user protection plan document (Subparagraph 3) contains objection procedures, error correction procedures, and relief methods for damage. For services that screen loans or filter job applicants using high-impact AI, it must set forth where and how a rejected user may file an objection, and through what procedures and within how many business days the company will review it.

The record of human management and supervision (Subparagraph 4) includes the name and contact information of the person in charge of managing and supervising the relevant high-impact AI (Article 27(1)(4) of the Enforcement Decree), supervision logs, and instances where human intervention actually occurred. If there is a record of a human intervening to modify a result when the AI made an anomalous determination, that serves as evidence of compliance

under Subparagraph 4.

The record regarding the preparation and retention of documents itself (Subparagraph 5) is a metadata document managing the preparation date, revision history, retention location, and access permission settings of the above four documents. A document is needed to prove the very fact that the documents are being retained.

Matters deliberated and resolved by the Committee (Subparagraph 6) apply when the National AI Strategy Committee has separately deliberated and resolved matters to secure safety and reliability of high-impact AI. As of July 2026, no specific resolution based on this provision has been made public; however, if the Committee determines industry-specific safety standards or additional requirements for specific technologies in the future, implementation records based on those resolutions will also be subject to retention.

3. Posting and Retention Are Different

A distinction must be made between Paragraph 1 and Paragraph 2 of Article 27 of the Enforcement Decree. Paragraph 1 concerns "posting." This means making information publicly available by posting it at offices, business premises, or on internet homepages. Paragraph 2 concerns "retention." This means maintaining internal

records.

Items subject to posting are limited to four: the major contents of the risk management plan, the major contents of the explanation plan, the user protection plan, and the name and contact information of the person responsible for management and supervision. Because the law specifies "major contents," there is no obligation to post the full text.

Matters that constitute trade secrets under Article 2, Subparagraph 2 of the Unfair Competition Prevention and Trade Secret Protection Act may be excluded from posting. Information such as detailed structures of model architectures, proprietary training data processing techniques, and weight parameters can be claimed as trade secrets. However, no such exception exists for the retention obligation. Even if information is a trade secret, documents must be internally prepared and retained for five years.

Items subject to retention encompass the entirety of Article 34(1)—that is, all documents serving as evidence of implementation for all six measures. The structure is one where posting selectively discloses while retention records everything without omission.

4. Format and Requirements for Electronic Documents

As long as the Enforcement Decree explicitly states "(including retention through electronic methods)," there is no practical reason to insist on paper documents. The problem is that the law has not set specific standards for the format of electronic documents.

The MSIT's Guidelines on the Obligations of High-Impact AI Business Operators do not prescribe a mandatory document format. Any form—whether PDF, Word files, spreadsheets, internal wikis, or Document Management Systems (DMS)—is acceptable as long as it remains in a state where "the contents of the measures can be verified."

Three formatting principles are recommended in practice.

First is searchability. When a fact-finding investigation is conducted five years later, it must be possible to immediately locate the documents from that point in time. Retaining documents solely as scanned images makes searching difficult. OCR-processed PDFs or structured databases are preferable.

Second is proof of integrity. It must be possible to prove that documents have not been tampered with after the fact. Systems that

automatically record creation timestamps, version control tools that maintain modification histories, and hash logging fall into this category. Some companies even adopt blockchain-based timestamps, but this is not a mandatory legal requirement.

Third is access control. Documents constituting trade secrets must have restricted access permissions to maintain trade secret protection requirements (non-public nature, economic utility, and secrecy management effort). Failing to maintain access logs may make it difficult to assert trade secret status later.

5. Practical Implementation of Retention Systems

While it is clear that documents must be prepared, the law is silent on where and how they should be stored. In practice, three levels of approach are possible.

The minimum level is a method of creating folder structures in existing internal file servers or cloud storage (such as Google Drive or Microsoft SharePoint) for retention. It incurs almost no cost. Defining document naming conventions (e.g., "Risk_Management_Plan_v2.1_20260315"), saving a new version with each edit, and setting access permissions will fulfill the minimum legal requirements. The problem is that after five years,

folders become scattered, and if the person in charge leaves the company, no one knows what is located where.

The intermediate level is a method of using an enterprise Document Management System (DMS) or electronic approval system. Products such as Hancom Docs by Hancom, Cloudium by CyberDym, and Microsoft Purview fall into this category. Versions are automatically assigned upon document registration, modification histories are recorded, and configuring retention policies prevents accidental deletion before five years elapse.

The highest level is a method of adopting an integrated GRC (Governance, Risk, Compliance) platform. Platforms such as ServiceNow GRC, Archer IRM, and OneTrust manage risk registers, control items, audit trails, and report generation within a single unified framework. If deployment histories, monitoring logs, and human intervention records of AI models are automatically collected, the burden of document preparation is reduced. In return, adoption costs appear to range from tens of millions to hundreds of millions of won.

Whichever level is chosen, the key is whether documents can be produced in front of an investigator five years later. Having a fancy

system is not what matters. What matters is that records remain complete, creation dates can be proven, and any intermediate edits can be tracked to show what changed.

6. The Boundary Between AI Logs and Documents

AI systems generate massive amounts of logs during operation. Inference request logs, input data, output results, response times, and error codes accumulate on a second-by-second basis. This raises the question of whether these logs themselves constitute the "documents" under Article 34.

The documents required by law are records that make it possible to "verify the contents of the measures." It must be possible to verify that a risk management plan was established, an explanation plan was implemented, and human supervision was conducted. Raw logs are traces of system operations, not the "contents of measures" in and of themselves.

However, there are cases where proving measures without raw logs is difficult. To prove "human management and supervision" (Subparagraph 4), records showing that a human reviewed and modified an AI determination on a specific date are required. These records are extracted from operational logs. To claim that "risk

mitigation measures were implemented," histories showing that a monitoring system detected anomalies and generated alerts serve as supporting evidence.

To summarize: there is no obligation to retain all raw logs for five years. However, it is safe to extract logs necessary to prove the implementation of measures and include them in documents or retain them separately. Safety check logs, human intervention logs, and anomaly response logs should preferably be included in retention targets. While there is no need to accumulate routine inference request logs for five years, in high-impact domains such as loan screening or hiring, retaining records of the basis for individual determinations is prudent in preparation for fact-finding investigations.

7. What Happens When a Fact-Finding Investigation Begins

Article 40 grants two powers to the Minister of Science and ICT: requesting data submission and conducting on-site investigations.

A request for data submission is a written request asking a business operator to submit relevant materials. It is triggered when a violation of the obligations of a high-impact AI business operator (Article 34(1)) is discovered or suspected, or when a report or complaint is filed.

Violations of documentation obligations are also included as grounds for fact-finding investigations. If documents are missing, it may be presumed that measures were not implemented.

An on-site investigation involves designated public officials entering offices or business premises to inspect account books, documents, and other materials or items. The procedures follow the Framework Act on Administrative Investigations. Written notification seven days prior to the investigation is the general rule, inspectors must present identification cards, and business operators may request the protection of trade secrets.

If a violation is recognized as a result of the investigation, the Minister of Science and ICT may order measures necessary to cease or correct the violating act (Article 40(3)). Failure to comply with a cessation order or a corrective order shall result in an administrative fine not exceeding 30 million won (Article 43(1)(3)).

Practical preparation requirements are clear. Internal targets should be set to submit documents within 30 days when requested by an investigator, and document locations and access methods should be organized in advance. Even if the person in charge resigns, a successor must be able to locate the files. If documents are scattered

across multiple departments, consolidating them alone could take weeks. Designating an AI compliance manager who manages the inventory and location of all documents is a fundamental preparatory step.

8. Technical Documentation under the EU AI Act: 10 Years and Conformity Assessment

The European Union AI Act imposes far more detailed and onerous requirements on technical documentation for high-risk AI systems.

Article 11 requires technical documentation to be drawn up before high-risk AI systems are placed on the market or put into service. The purpose of the documentation is to demonstrate that the relevant system complies with the requirements set by law, and it must contain sufficient information for national competent authorities and notified bodies to assess such compliance.

The contents of technical documentation are specified in Annex IV. They include a general description of the system (intended purpose, developer information, version), detailed descriptions of system elements (architecture, algorithms, data processing procedures), information on training data (data collection methods, preprocessing, labeling, statistical characteristics), descriptions of

monitoring, functioning, and control, information on implementation of the risk management system, records of conformity assessment procedures, and changes made throughout the lifecycle.

The retention period is 10 years. Article 18(1) provides that providers of high-risk AI systems shall keep the technical documentation for 10 years after the system has been placed on the market or put into service. General-purpose AI (GPAI) models must also draw up technical documentation and keep it up to date pursuant to Article 53. However, the 10-year retention of GPAI technical documentation is not a statutory obligation under the text of the law, but a voluntary compliance standard established by the GPAI Code of Practice published in July 2025.

Five years in South Korea, ten years in the EU. Looking at numbers alone, it is a two-fold difference. However, the real difference lies not in the duration, but in the purpose of the documents. Technical documentation in the EU serves as evidentiary material for conformity assessment. To place high-risk AI on the market, it must pass a prior conformity assessment, and supervisory authorities can request technical documentation at any time to re-verify compliance. Korea's documents serve as evidence for ex-post fact-finding

investigations. There is no obligation to submit them in advance; supervisory authorities request them when there are suspected violations.

For business operators operating in both the EU and South Korea, the practical implications are clear. Preparing technical documentation to EU standards will automatically satisfy Korean requirements. Conversely, meeting only Korean requirements will necessitate building documentation from scratch upon entering the EU. If targeting both markets, it is more efficient from the start to prepare documents at the EU level while setting the retention period to 10 years.

9. China's Filing and Record-Keeping Obligations

China imposes documentation obligations centered on an algorithm filing (备案) system. According to Article 17 of the Interim Measures for the Management of Generative Artificial Intelligence Services, generative AI services with public opinion attributes or social mobilization capabilities must undergo a security assessment and file their algorithms with the Cyberspace Administration of China (CAC) before providing services to the public. Even after filing, if significant changes occur in the principles or functions of the algorithm, filing information must be updated within 10 business

days.

The Measures for the Labeling of AI-Generated and Synthetic Content, implemented in September 2025, require retaining implementation records of dual labeling (explicit labels and metadata or digital watermarks). Mandatory national technical standards (GB 45438-2025) specify the technical specifications for labeling methods.

The difference from South Korea lies in whether prior filing is required. In China, if a service falls under filing obligations, documents must be submitted to the authorities before launch. South Korea requires documents to be retained internally after launch and produced when requested by authorities. The burden of document preparation is similar, but the timing of submission differs.

Comparing the three countries, the intensity of documentation obligations is clear. The EU is the heaviest: prior conformity assessment grounds, 10-year retention, and detailed Annex IV lists. China comes next: prior filing, updates upon major changes, and compliance with national technical standards. South Korea is the lightest: ex-post retention, 5 years, and flexible formatting.

10. Dual Obligations and Mitigation of Duplicate Regulation

A significant number of high-impact AI business operators already bear documentation obligations under other laws. Financial companies retain electronic data for five years or more under the Electronic Financial Transactions Act and the Use and Protection of Credit Information Act. Medical device manufacturers retain design records and quality records under the Medical Devices Act. Business operators processing personal information retain processing records under the Personal Information Protection Act.

Article 34(3) of the AI Basic Act and Article 27(5) (Attached Table 1) of the Enforcement Decree mitigate this duplication. Where measures corresponding to each subparagraph of Article 34(1) have been implemented pursuant to other statutes or regulations, they are deemed to have fulfilled the measures under the AI Basic Act.

For example, if a financial company operating a loan screening AI complies with the obligation to explain under Article 35-2 of the Credit Information Act and establishes procedures for fulfilling the obligation to explain under Article 36-2(1) of the same Act, it is recognized as having implemented Article 34(1)(2) (explanation plan) of the AI Basic Act. The user protection plan (Subparagraph 3) may be recognized if all obligations under Article 10 of the Financial

Consumer Protection Act have been fulfilled. Conversely, loan screening is not included among targets recognized for risk management plans (Subparagraph 1), meaning it must be implemented separately. Which measures under which laws are recognized are organized by law in Attached Table 1 of the Enforcement Decree.

However, duplication mitigation does not equate to an exemption. If measures under other laws do not cover all subparagraphs of Article 34, the remaining obligations must be implemented separately and documented. Even if a financial company obtains recognition for an explanation plan under the Credit Information Act, records for risk management plans under Subparagraph 1, "human management and supervision" under Subparagraph 4, and "document preparation and retention" itself under Subparagraph 5 remain unique requirements of the AI Basic Act.

11. The Reality of Document Management Automation

Tools that automatically generate and manage AI compliance documentation are available on the market.

Model Cards and Datasheets are standardized document templates proposed by research teams at Google and Microsoft. They

summarize on a single sheet the model's intended purpose, training data, performance metrics, bias evaluation results, intended scope of use, and out-of-scope uses. Model hubs like Hugging Face provide features to automatically generate Model Cards.

MLOps pipeline tools (such as MLflow, Weights & Biases, and Neptune) automatically record model training hyperparameters, dataset versions, performance metrics, and deployment histories. These records can serve as evidentiary documentation for Article 34(1)(2) (explanation plan). Because experiment records are automatically logged every time training is run, there is no need to reconstruct "what data was used and how it was trained" after the fact.

Privacy and AI governance platforms such as OneTrust and BigID centrally manage AI inventories (lists of AI systems operated by an organization), risk assessments, and impact assessment results. Many of their features were developed for EU AI Act compliance, so slight customization is required to adapt them to Korean legal requirements.

Open-source tools are also available. Representative examples include Google's SAIF (Secure AI Framework) checklist and the NIST

AI RMF (Risk Management Framework) self-assessment tools. They can be initiated without financial burden, but human resources are required for implementation and maintenance.

Whether tools are adopted or management is conducted via Excel, what determines whether documentation satisfies legal requirements is the content, not the tool. Automated tools will not fill in the blanks for you. Identifying and evaluating risks requires human judgment, and tools serve merely as the means to record that judgment.

12. Practical Responses for Small and Medium Enterprises (SMEs)

Telling an AI startup with 20 employees to adopt a GRC platform is unrealistic. Let us consider methods for SMEs to meet legal requirements with minimal cost and workforce.

First, business operators must check whether their AI falls under high-impact classification. If it does not fall under any of the 11 domains specified in Article 2, Subparagraph 4 (such as energy, healthcare, hiring, loan screening, transportation, public services, and student evaluation), the documentation obligations under Article 34 do not apply in the first place. If high-impact status is uncertain, a request for confirmation may be submitted to the MSIT (Article 33(1)).

Requests for confirmation are free of charge, and a response is provided within 30 days.

If it qualifies as high-impact, a minimum document set must be created. Create a risk register in a single Google Spreadsheet and log identification dates, severity, likelihood of occurrence, mitigation measures, and mitigation status for each risk item. Outline the explanation plan and user protection plan in a single Word document. Summarize the sources, scale, and preprocessing methods of training data. Designate a manager responsible for management and supervision, and post their name and contact information on the homepage. Store these documents in cloud storage within date-based folders and save as new files upon each revision to maintain version control.

The AI Basic Act Help Desk (swai.or.kr/AI_act_helpdesk) answers such practical questions free of charge. In just 10 weeks of enforcement, 552 consultation requests were received, and anonymous consulting is also available. In-depth legal reviews take up to 14 days. The Guidelines on the Obligations of High-Impact AI Business Operators are also publicly available free of charge.

If you are a deployer business operator using AI supplied by a developer business operator, you should check Article 27(3) of the Enforcement Decree. If the developer business operator has already implemented the measures under Article 34(1)(1) through (3), and the deployer business operator has not significantly altered the original purpose or scope of use of the relevant AI, the deployer business operator is recognized as having complied to that extent. Even in such cases, it is safe to retain documentation verifying the fact of compliance itself (e.g., a copy of the developer's risk management implementation confirmation). Article 27(4) of the Enforcement Decree provides that deployer business operators may request developer business operators to provide materials necessary to comply with Article 34.

13. What Happens If There Are No Documents?

The administrative fine provisions of the AI Basic Act (Article 43) do not directly enumerate document retention violations as grounds for administrative fines. Grounds for administrative fines are limited to three: violation of transparency obligations (Article 31(1)), failure to designate a domestic representative (Article 36(1)), and non-compliance with a cessation order or corrective order (Article 40(3)).

This does not mean document retention can be ignored. The sequence of events is as follows: if a lack of documentation is discovered during a fact-finding investigation under Article 40, a corrective order may be issued for violation of Article 34(1). Failure to comply with the corrective order results in an administrative fine. Because the absence of documents serves as evidence that measures were not implemented, document retention is effectively a prerequisite for fulfilling all obligations.

The government announced that it will operate a grace period (guidance period) of at least one year following the law's enforcement. Under this policy, administrative guidance and corrective recommendations will take precedence for omissions of transparency obligations or documentation deficiencies. Serious matters—such as deepfake criminal exploitation, loss of human life, or large-scale human rights violations—are exceptions.

Once the grace period ends, the situation changes. Without documents, there is no way to prove whether risk management was conducted, an explanation plan was established, or human supervision was exercised. In a fact-finding investigation, asserting "we complied but did not record it" is unlikely to be accepted.

Without records, there is no evidence to support compliance, leaving the outcome of the investigation to lean disadvantageously against the business operator.

14. Domestic Representative's Duty to Inspect Documents

Where an overseas AI business operator operates in South Korea and designates a domestic representative, the domestic representative also has a document-related role. Article 36(1)(3) includes within the scope of duties of a domestic representative "support necessary for the implementation of measures to secure safety and reliability pursuant to each subparagraph of Article 34(1)," explicitly adding in parentheses, "(including inspections regarding the recency and accuracy of documents under Subparagraph 5 of the same paragraph)."

The domestic representative must inspect whether the documents prepared by the overseas headquarters are up to date and accurate. This is not a mere formal verification, but a substantive review. It must verify whether the risk management plan reflects the circumstances of the Korean market, whether the explanation plan is appropriate for Korean-language users, and whether the management and supervision manager's information is currently valid.

From the perspective of overseas business operators, because documents must be shared with the domestic representative, a Non-Disclosure Agreement (NDA) must be established beforehand. Non-disclosure obligations must be clearly defined when sharing technical information that constitutes trade secrets with the domestic representative.

15. Documents Are Evidence

Let us return to the fintech company in Gangnam, Seoul, in April 2026. What if, when the investigator requested the risk management plan document and implementation evidence records, the company had produced systematic documentation? If loan denial bias risk had been identified in the risk register, weight adjustments of specific variables and periodic bias audits recorded as mitigation measures, and recent audit results attached, the investigator could have judged that "this company is complying with Article 34." The company could also have explained through documents that the cause of the complaint was not systemic bias, but individual differences (income structure, debt ratio, etc.).

Document management is both a defensive measure in preparation for supervision and a tool for business operators themselves to inspect their own systems. In the process of articulating risks into

documentation, loopholes are uncovered. While drafting explanation plans, companies may face the reality that "we ourselves do not know why this model produces such results." That realization is the beginning of improvement.

A retention period of five years means that even if technology changes, personnel change, and business models change in the interim, records must remain intact. In an era where AI increasingly makes determinations on our behalf, recording and retaining the rationale for those determinations has become a legal obligation. What the law ultimately asks comes down to one thing: Why did your AI reach such a determination? Can you produce the answer to that question five years from now?

Part 8. Entering Overseas Markets and How to Read International Competitors

Chapter 24. South Korean Law Side-by-Side with the EU AI Act

On February 2, 2025, social scoring AI systems were banned across the European Union. This is technology through which governments score citizens' online behavior and social reputation, restricting access to public services based on those scores. It was the day Article 5 of the EU AI Act came into effect. At the exact same time, the CEO of an AI startup in Seoul asked their legal team what preparation was needed for their recruitment interview analysis service to enter the European market. It took three weeks for an answer to come back. The response indicated that while they believed complying with South Korean law would be sufficient, they might have to redesign their entire business model in the face of EU law.

While South Korea's AI Basic Act and the EU AI Act regulate the same subjects during the same period, their underlying legislative philosophies differ from the start. This divergence spans six layers: terminology choice, classification schemes, prior control mechanisms, levels of sanction, prohibited categories, and governance structures.

1. Risk vs. Impact: The Classification Philosophies of the Two Laws

The EU looks at risk. South Korea looks at impact.

The EU AI Act classifies AI systems into four tiers. Unacceptable Risk is banned outright. High-Risk systems must undergo a prior conformity assessment before entering the market. Limited Risk comes with transparency obligations. Minimal Risk operates freely without separate regulations. From the top of the pyramid to the bottom, it is a structure in which obligations vary according to the severity of risk.

South Korea's AI Basic Act establishes a single category known as 'high-impact AI.' It targets AI systems that significantly impact or carry the risk of posing a threat to human life, bodily safety, and fundamental rights (Article 2, Item 4). Ten domains—including energy supply, healthcare, recruitment, credit screening, transportation, public services, and student evaluation—are directly enumerated in Article 2, Item 4 of the Act, accompanied by a delegating clause allowing additional domains to be added by the Enforcement Decree. AI that is not high-impact is not subject to separate regulations, with the exception of transparency obligations and safety assurance obligations (Article 32) attached to systems meeting large-scale compute criteria.

Politics played a role in the choice of terminology. During the Science, ICT, Future Planning, Broadcasting and Communications Committee review, the expression 'high-risk' also appeared in early bill drafts. Assemblywoman Lee Hae-min opposed it, remarking: "The reason the global community is moving toward 'impact' is that impact can encompass both riskiness and scope." The intent was not to stigmatize the technology itself. While the EU focuses on potential harm through the word 'high-risk,' South Korea shifted its gaze to social influence through the word 'high-impact.'

This choice comes at a price. Research Fellow Jung Won-jun of the Korea Legislation Research Institute pointed out in a September 2025 report that confusion arises because 'impact' and 'risk' are regulated side by side within the statutory text. The definition clause of high-impact AI itself contains the phrase 'concern of causing risk.' Having declared a focus on impact, the criteria actually reintroduce risk back into the evaluation. While the EU breaks risk down into four detailed tiers and assigns obligations tailored to each tier, South Korea imposes uniform obligations on all high-impact AI within a single category. This is why the Korea Legislation Research Institute and academic papers (Baek Su-won, European Constitutional Law Research Vol. 48) repeatedly raise the limitations of this single-tier

structure.

2. Prior Conformity Assessment vs. Autonomous Impact Assessment

Under the EU AI Act, a high-risk AI system must pass a conformity assessment to be placed on the market. If it falls under product safety regulations (Annex I), an assessment by a third-party notified body is required. Standalone high-risk AI under Annex III performs a self-conformity assessment but must prepare technical documentation, a quality management system, and an EU declaration of conformity. Failing the conformity assessment prevents obtaining the CE marking, and without the CE marking, it cannot be sold in the EU market.

South Korean law has no such prior gateway. Obligations imposed on high-impact AI business operators are contained in five items under a single article, Article 34, Paragraph 1. These are: establishment and operation of risk management measures (Item 1), establishment and implementation of explanation measures (Item 2), establishment and operation of user protection measures (Item 3), human oversight and control (Item 4), and drafting and retention of documents (Item 5). The five-year retention period is stipulated in Article 27, Paragraph 2 of the Enforcement Decree. Added to this is the fundamental rights impact assessment (Article 35, Paragraph 1),

which is not a strict mandatory obligation but a 'duty of autonomous effort.' Although the Ministry of Science and ICT (MSIT) guidelines outline three stages of impact assessment (preliminary preparation, main assessment, and post-action measures), the guidelines themselves carry no legally binding force.

Being autonomous does not mean it lacks all enforcement power. Business operators that carry out impact assessments are given preferential consideration in public procurement. Article 16, Paragraph 3 of the Act stipulates that when state agencies and related entities purchase products or services necessary for their operations or commission services, they must give preferential consideration to AI products and services verified by the Minister of Science and ICT, with criteria and procedures for such verification determined by administrative notice (Article 15, Paragraphs 4 and 5 of the Enforcement Decree). Article 35, Paragraph 2 stipulates that when state agencies and related entities intend to use products or services utilizing high-impact AI, they must give preferential consideration to products and services that have conducted impact assessments. Although the statute does not mandate the impact assessment itself, market access conditions through government procurement act as a 'soft regulation' that de facto induces

compliance.

Placing the two frameworks side by side reveals distinct directions. The EU says, "Prove safety before coming to market." South Korea says, "Manage yourselves after entering the market, and show an assessment if you want to enter the government market." The EU approach favors consumer protection but may slow down the pace of innovation. The South Korean approach reduces the initial burden on business operators, but risks becoming a reactionary measure applied only after problems occur. Neither approach is flawless. The question is which type of failure one is willing to accept.

3. Provider/Deployer vs. Development Business Operator / User Business Operator

The EU AI Act tightly categorizes actors participating in the AI value chain. A Provider is one who develops an AI system or general-purpose AI model and places it on the market under their own name. Whether developed directly or outsourced and launched under their own trademark, they are a Provider. A Deployer is one who uses an AI system under their authority in a professional context. If a hospital adopts an AI medical image reader to treat patients, that hospital is a Deployer. Added to these are separately defined roles including Importer, Distributor, Authorized Representative, and

Product Manufacturer. At least seven roles appear in the statutory text.

South Korea's AI Basic Act uses only two categories. A development business operator (Article 2, Item 7, Subitem (a)) is one who develops and provides artificial intelligence. A user business operator (Subitem (b) of the same Item) is one who provides products or services using AI provided by a development business operator. In addition, there is the user (Article 2, Item 8), defined as the entity receiving AI products or services—namely, the end consumer.

Practically speaking, the decisive difference lies between Deployer and user business operator. In the EU, when a hospital uses an AI medical image reader, it is a Deployer. It must perform a fundamental rights impact assessment, fulfill transparency obligations, and establish human oversight systems. In South Korea, if the same hospital uses the exact same device solely for its own diagnostic duties, it remains classified as a 'user,' incurring no regulatory obligations. However, if structured as providing medical services to patients using AI supplied by a developer, it may fall under user business operator (Article 2, Item 7, Subitem (b)), in which case transparency obligations and high-impact business operator

responsibilities apply.

This difference presents a practical barrier for South Korean business operators entering the EU. In South Korea, the compliance burden on institutions adopting AI is dramatically low. In the EU, institutions adopting AI also bear substantial obligations as Deployers. For a South Korean AI enterprise to supply systems to European hospitals or schools, it must simultaneously provide technical documentation, user manuals, and monitoring tools so that those hospitals or schools can fulfill their obligations as Deployers. Companies that previously complied solely with South Korean law suddenly find themselves forced to establish Deployer support systems when tapping into the EU market.

One more thing. The EU's definition of Provider includes the phrase "under its own name or trademark." White-label operators who take external models and offer services under their own brand are also Providers. South Korean law lacks an explicit counterpart to this provision. Whether an entity that fine-tunes external models to deliver new services qualifies as a development business operator or a user business operator remains an area awaiting official interpretation from the Ministry of Science and ICT.

4. A Thousand-Fold Gap in Fines

Let us set the numbers side by side.

The maximum fine under the EU AI Act is up to 35 million EUR or 7% of total worldwide annual turnover, whichever is higher. Based on the exchange rate in July 2026, this translates to approximately 61 billion KRW. Violating obligations for high-risk AI incurs fines of up to 15 million EUR or 3% of global turnover. The highest tier of sanctions applies to prohibited practice violations.

The maximum administrative fine under South Korea's AI Basic Act is 30 million KRW. It is imposed in three specific instances (Article 43, Paragraph 1): failure to provide prior notification of AI usage (Article 31, Paragraph 1), failure to designate a domestic representative (Article 36, Paragraph 1), and failure to comply with suspension orders or corrective orders (Article 40, Paragraph 3). Failing to fulfill documentation obligations does not trigger an immediate administrative fine; it leads to an administrative fine only if an operator fails to comply after receiving a corrective order.

Approximately 61 billion KRW versus 30 million KRW. A gap exceeding a thousand times. Criticism existed even before the law took effect that 30 million KRW in South Korea is insufficient to

motivate large corporations with annual revenues of hundreds of billions of KRW to ensure compliance. When the level of administrative fines became controversial during the committee review, the rationale for keeping the amount low was that domestic AI industries were in their early growth stage and excessive penalties would stifle innovation. Accounting for grace periods as well, administrative fines will not even be levied for omissions of transparency obligations prior to January 2027.

There is no guarantee that the EU's heavy fines are strictly effective. Since GDPR took effect, fines amounting to tens of millions of euros have been levied, yet companies frequently file legal challenges and litigate for years. Concluding that South Korea's low administrative fines are necessarily ineffective is equally premature. Reputational damage or exclusion from public procurement may act as stronger compliance incentives than administrative fines.

Nonetheless, a thousand-fold disparity numerically illustrates the difference in enforcement resolve between the two laws. The EU enshrined in statutory text its determination to penalize regulatory violations to a degree that threatens business survival. South Korea chose a strategy of eliciting compliance through market incentives

rather than punitive measures.

5. Prohibited Practices: What Is Absent in South Korea

Article 5 of the EU AI Act bans specific AI practices outright. It is worth examining the complete list of prohibitions.

Social scoring is prohibited. Whether by public entities or the private sector, scoring social behavior via AI to impose disadvantages is disallowed. Subliminal manipulation is also banned, targeting AI that influences behavior in ways individuals cannot perceive or AI that exploits vulnerabilities. Real-time remote biometric identification in law enforcement is prohibited in principle, with only extremely narrow exceptions such as anti-terrorism. Emotion recognition AI in workplaces and educational institutions falls within the banned scope as well, subject to exceptions for medical or safety purposes. With the 2026 Digital Omnibus amendments, non-consensual intimate image generation AI and child sexual abuse material generation AI were added (applicable from December 2, 2026).

South Korea's AI Basic Act contains no such prohibited categories. There are no statutory provisions blocking specific types of AI at their source. While it imposes risk management and transparency

obligations on high-impact AI business operators, it nowhere declares that AI for specific purposes cannot be developed or operated at all.

Whether this absence reflects a weakness or flexibility in South Korean law depends on one's perspective. EU civil society views prohibition clauses as the baseline for fundamental rights protection, arguing that a society permitting social scoring is fundamentally different from one that bans it. Portions of South Korean industry counter that in a rapidly evolving technological landscape, pre-emptively banning specific use cases risks blocking unanticipated beneficial applications as well.

From a practical standpoint, a key fact remains: if a South Korean business operator wishes to offer services in the EU, it must comply with the EU's prohibition list in its entirety. The absence of prohibitions in South Korean law does not mean a South Korean enterprise can operate social scoring AI in the EU market. Article 5 applies to all Providers and Deployers supplying output within the EU territory.

6. General-Purpose AI (GPAI) vs. Generative AI

The EU AI Act introduces the concept of General-Purpose AI (GPAI). A GPAI model (Article 3, Point 63) is a model "trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks." A GPAI system (Article 3, Point 66) is a system built upon such a model that can be used for a variety of purposes. Added to this is an upper tier called 'systemic risk.' If cumulative training compute reaches 10^{25} FLOPs or higher, it is presumed to be a GPAI with systemic risk, incurring stricter obligations (risk assessments, incident reporting, and securing cybersecurity).

South Korean law has no concept corresponding to GPAI. Instead, it defines 'generative artificial intelligence' (Article 2, Item 5) as "an artificial intelligence system that generates output such as text, audio, images, or video by imitating the structure and characteristics of input data." Article 32 of the Act and Article 24 of the Enforcement Decree specify the targets for safety assurance obligations—a category commonly referred to as 'high-performance artificial intelligence.' This applies to systems that satisfy all three requirements: cumulative training compute of 10^{26} FLOPs or more, application of state-of-the-art AI technology, and a level of risk

carrying concerns of widespread and significant impact on human life, bodily safety, and fundamental rights.

There are two key differences. Regarding compute thresholds, the EU uses 10^{25} while South Korea uses 10^{26} —a tenfold difference. A model subject to regulation under EU standards may remain outside the regulatory scope under South Korean standards. As of July 2026, virtually no AI models developed domestically exceed 10^{26} FLOPs. This effectively narrows the scope of regulation in practice.

The conceptual scope also differs. The EU's GPAI is not limited to content generation. Any model possessing 'general-purpose capabilities'—such as classification, reasoning, tool use, or code writing—qualifies as GPAI. South Korea's 'generative AI' focuses specifically on systems that 'generate' output. Whether an AI model that does not generate content but possesses general-purpose reasoning capabilities falls under the definition of generative AI in South Korean law remains unclear.

7. EU AI Office vs. National AI Strategy Committee

The central body overseeing the enforcement of the EU AI Act is the European AI Office. Established within the European Commission on February 21, 2024, it comprises over 125 staff members across six

units: AI and Robotics Excellence, Regulation and Compliance, AI Safety, AI Innovation and Policy Coordination, AI for Societal Good, and AI in Health and Life Sciences. It holds direct enforcement authority over GPAI models, with supervisory authority added over AI systems integrated into Very Large Online Platforms (VLOPs) and Very Large Online Search Engines (VLOSEs) under the Digital Omnibus revisions. In addition, each Member State maintains its own National Competent Authority to conduct market surveillance of high-risk AI systems.

South Korea's governance structure differs. The National AI Strategy Committee serves as the supreme policy coordination body. Chaired by the President, it operates with up to 60 members across eight subcommittees. Authority over policy coordination and performance management was expanded under the 2026 revised Act. Practical enforcement is handled by the Ministry of Science and ICT (MSIT). The National Information Society Agency (NIA), designated as the AI Policy Center, supports policy development and master plans, while the AI Safety Institute (affiliated with ETRI) handles safety evaluations and safety technology research.

The difference lies in enforcement power. The European AI Office possesses direct powers of investigation and fine imposition over GPAI Providers. Upon confirming a violation, it can independently issue corrective orders and impose sanctions. South Korea's Strategy Committee is a policy coordination body, not an enforcement agency. While MSIT handles fact-finding investigations and levies administrative fines, administrative guidance and corrective recommendations take priority during the grace period. Although the Council of Chief AI Officers (CAIOs) across 48 central administrative agencies meets monthly, this functions as an internal government coordination mechanism rather than an enforcement body directed at business operators.

8. Additional Compliance Requirements for South Korean Business Operators in the EU Market

Business operators that previously complied solely with South Korean law must fulfill additional obligations not present in South Korean law when entering the EU market. Let us examine the detailed list.

Reviewing prohibited practices comes first. Operators must verify that their AI does not fall under the prohibition list in Article 5 of the EU AI Act. Because South Korean law lacks prohibited categories, AI

features that are lawful in South Korea may be banned in the EU. AI analyzing applicant emotions during recruitment processes or AI recognizing student emotional states in schools serve as representative examples.

High-risk AI classification must be re-conducted. The EU's high-risk determination consists of a dual pathway (the product safety pathway and the Annex III pathway) alongside a derogation mechanism (Article 6, Paragraph 3). While South Korea's 11 high-impact domains and the EU's Annex III domains overlap substantially, they are not entirely identical. The EU lists biometrics, critical infrastructure, education, employment, access to essential services, law enforcement, migration and border control, administration of justice, and democratic processes. South Korean law contains no separate entries for migration management or democratic processes.

Conformity assessments must be passed. While an autonomous impact assessment suffices in South Korea, entering the EU market requires preparing technical documentation, a quality management system, an EU declaration of conformity, and obtaining the CE marking. Third-party certification may also be required in certain

cases.

Deployer support systems must be established. Most institutions adopting AI in South Korea fall outside regulation as mere 'users,' but identical institutions in the EU act as Deployers. When South Korean AI companies supply systems to European clients, they must provide user manuals, logging systems, and human oversight tools to enable those clients to fulfill their Deployer obligations.

GPAI obligations must be verified. Even if a model does not fall into South Korea's 'generative AI' or 'high-performance AI' categories, qualifying as GPAI under EU standards triggers obligations for technical documentation, publishing training data summaries, and complying with copyright policies. If classified as a GPAI with systemic risk, risk assessments, incident reporting, and cybersecurity measures are additionally required.

An Authorized Representative must be designated. Providers without an establishment in the EU must appoint an Authorized Representative within the EU territory. While similar to South Korea's domestic representative system, the role and scope of liability for an EU representative are broader.

The level of transparency obligations differs. South Korean law requires prior notification of AI operation and labeling of AI-generated output. In addition to these, the EU requires separate disclosures that an AI system is interacting with humans, that an emotion recognition system is active, and that content is deepfake media (Article 50).

Attention must also be paid to deadlines. Under Digital Omnibus amendments, application dates for high-risk AI obligations were postponed to December 2, 2027 for Annex III standalone systems and August 2, 2028 for Annex I product-embedded systems. Although this grace period buys time to prepare, prohibitions (enforced February 2025), GPAI obligations (enforced August 2025), and transparency obligations (enforced August 2026) are already in effect or will take effect soon.

9. To Business Operators Standing Between the Two Laws

South Korea's AI Basic Act grants business operators broad autonomy. The EU AI Act imposes intricate obligations on business operators. That South Korean law is flexible does not mean EU law is inherently correct, nor does the strictness of EU law imply that South Korean law is flawed. The two laws reflect distinct societal consensus. South Korea determined that AI industry growth was the immediate

priority, whereas the EU judged that protecting citizens' fundamental rights takes precedence over industrial promotion.

The question remaining for business operators is not philosophical debate, but practical implementation. If there is even a slight possibility of your AI service entering the EU market, compliance with South Korean law alone is insufficient. You must assess right now whether your service qualifies as high-risk under EU standards, whether it violates prohibited practices, whether it can pass a conformity assessment, and what documentation must be provided to Deployers.

South Korea's AI Basic Act is the world's first comprehensive AI legislation to enter into full effect. The EU AI Act's high-risk obligations apply starting December 2027. Since South Korea started earlier, South Korean business operators have time to bridge the gap between the two laws. How much time remains depends on the business model. The moment you send your first proposal to a European client, the clock is already ticking.

Chapter 25. Practical Implications of the 2026 EU Omnibus Revision for Korean Business Operators

On November 19, 2025, the European Commission announced the Digital Omnibus Bill (COM(2025) 836 final), which adjusts the timing for applying the key obligations of the AI Act. It contained provisions postponing the application dates of key obligations of a law that had been promulgated only 1 year and 4 months prior. In May 2026, the European Parliament and the Council reached a provisional agreement, which was confirmed by the member states, leaving only publication in the Official Journal remaining. The entire regulatory timetable for high-risk AI is changing. For business operators entering or preparing to enter the EU market from Korea, this presents both an opportunity and a risk to reset their preparation strategies.

1. Background of the Omnibus Revision: Why Was It Amended So Quickly?

The EU AI Act (Regulation 2024/1689) entered into force on August 1, 2024. It was designed so that most provisions would apply starting August 2, 2026. However, an amendment proposal emerged approximately 15 months after its entry into force.

Behind this was strong backlash from the European AI industry. European AI companies such as France's Mistral, Germany's Aleph Alpha, and Finland's Silo AI voiced concerns over the burden of

conformity assessment costs and the contraction of investment due to regulatory uncertainty. The logic was that European companies alone were being held back while non-European competitors (such as OpenAI in the US and Baidu in China) raced ahead under relatively lighter regulation.

The second term of the European Commission under President Ursula von der Leyen, launched in December 2024, put easing the regulatory burden as a key agenda item. The Omnibus revision is a concrete product of this agenda.

2. Postponed Obligations: What Has Been Pushed Back?

The core of the amendment proposal is the postponement of the application dates for obligations regarding high-risk AI systems by category.

Standalone high-risk AI systems have been pushed back from August 2, 2026, to December 2, 2027. This provides an additional grace period of approximately 16 months. AI systems in sectors listed in Annex III (such as biometrics, critical infrastructure, education, employment, and law enforcement) fall under this category.

AI systems embedded in product safety legislation have been postponed from August 2, 2027, to August 2, 2028. This applies when AI is incorporated into products subject to existing product safety regulations, such as medical devices, automobiles, aviation, and machinery. This category has a structure where existing product certification regimes and AI Act conformity assessments apply doubly; the postponement reflects the judgment that coordinating the two regimes requires more time.

The effect of this grace period on Korean business operators is clear. High-risk AI business operators preparing to enter the EU market have gained an additional 16 months to prepare for conformity assessments. Business operators exporting medical device AI now have breathing room until August 2028.

However, a grace period is not an exemption. Even during the grace period, the EU AI Office will issue guidelines for high-risk AI systems and carry out work on harmonised standards. The aforementioned dates represent deadlines meaning "by then at the latest," so if the Commission confirms that compliance support tools such as standards are in place, application could occur earlier. The grace period must not be viewed as a rest break; it must be used as

preparation time.

3. Non-Postponed Obligations: What Enters into Force on August 2, 2026?

Not everything has been pushed back. There are obligations already in effect, and obligations that will apply on August 2, 2026, according to the original schedule.

Obligations for General-Purpose AI (GPAI) models (Articles 51–55) have already been in effect since August 2, 2025. GPAI model providers bear the obligations to draw up technical documentation, establish policies to comply with copyright law, and publish a sufficiently detailed summary of training data. GPAI models with systemic risks are additionally subject to obligations including red-teaming tests, cybersecurity measures, and reporting serious incidents.

Transparency obligations (Article 50) will take effect on August 2, 2026, as scheduled. These include obligations to disclose that an AI system interacts with humans (chatbot notice), label AI-generated content (including deepfake labeling), and label AI-generated text when dealing with matters of public interest. However, through the Omnibus agreement, an additional 4-month grace period until December 2, 2026, was granted for the machine-readable labeling

(watermarking) requirements of existing systems.

Enforcement powers over GPAI will also be activated on August 2, 2026. The European Commission (specifically the AI Office within it) will directly supervise GPAI model providers and can impose administrative fines (up to 3% of total worldwide annual turnover or 15 million euros, whichever is higher) for violations.

Among Korean business operators, those providing chatbots or generative AI services to the EU market must prepare for transparency obligations by August 2, 2026. This point in time is only 12 days apart from the enforcement of the amended Korean AI Basic Act (July 21, 2026). Establishing a system that satisfies the transparency obligations of both laws simultaneously can reduce duplicate work.

4. Comparative Timeline with the Enforcement Dates of Korean Law

Placing the enforcement dates of the two laws side by side reveals a pattern.

February 2, 2025: Application of prohibited AI practices (Article 5) under the EU AI Act begins. It is already in effect.

August 2, 2025: Application of General-Purpose AI (GPAI) model obligations under the EU AI Act begins. Obligations for GPAI technical documentation, copyright policies, and training data summaries took effect from this date.

January 22, 2026: First-phase enforcement of the Korean AI Basic Act. Basic obligations (prior notification, output labeling) along with obligations for high-impact and high-performance AI take effect. However, implementation is expected to focus on guidance in the early stage.

July 21, 2026: Enforcement of the amended Korean law. Amended and newly established provisions, such as expanding accessibility and promoting investment, take effect.

August 2, 2026: Application of transparency obligations (Article 50) under the EU AI Act begins and GPAI administrative fine powers are activated. Chatbot notification and deepfake labeling.

2027 onwards: End of Korea's guidance-focused operation (expected). Full-scale enforcement begins.

December 2, 2027: Application of EU standalone high-risk AI obligations (post-Omnibus revision timeframe).

August 2, 2028: Application of EU product-embedded high-risk AI obligations.

If Korean business operators enter the EU market simultaneously, they can establish a roadmap to implement transparency obligations for both sides in the second half of 2026 and initiate preparations for EU high-risk obligations starting no later than the first half of 2027.

5. Practical Roadmap for Korean Business Operators Preparing for the EU Market

Specifically, what needs to be done and when?

Immediately: If you are a GPAI model provider, inspect the status of drawing up technical documentation, establishing copyright policies, and publishing training data summaries. As this obligation has been in effect since August 2, 2025, if it has not yet been fulfilled, it must be implemented right now.

First half of 2026: Identify transparency obligations under Article 50 of the EU AI Act. Prepare a comparison table mapping overlapping areas and additional requirements against the prior notification and

output labeling obligations of the Korean law. Design notification phrasing, labeling methods, and technical measures that satisfy both simultaneously.

Second half of 2026: Apply chatbot notifications and deepfake labeling to services targeting the EU market.

First half of 2027: Track the issuance status of EU harmonised standards. Once the AI quality management standards and risk management standards being developed by CEN-CENELEC are finalized, self-assess product conformity. Determine by this time whether your standalone high-risk AI system falls under the Annex III sectors, and if so, begin preparing for conformity assessment procedures (third-party certification or self-declaration of conformity).

Second half of 2027: If third-party certification is required, inquire about certification schedules with an EU-notified body (Notified Body) and proceed in earnest with conformity assessments. Because notified bodies have limited processing capacity, early action is necessary.

First half of 2028: If you have product-embedded AI, complete integrated documentation combining existing product certifications (CE marking, MDR, Machinery Regulation) and AI Act conformity assessments.

6. EU Regulatory Sandboxes and Korea's Regulatory Sandbox

Article 57 of the EU AI Act provides for AI regulatory sandboxes. This is a framework allowing member states to test AI systems in a controlled environment and verify in advance whether regulatory requirements are met.

Korea's regulatory sandbox consists of temporary authorizations based on Article 37 and demonstration special cases based on Article 38-2 of the Information and Communications Convergence Act (ICT Special Act). Although the AI Basic Act did not establish a separate sandbox framework, Article 19(3) stipulates that the government shall actively support the smooth implementation of these temporary authorizations and demonstration special cases for AI-converged products and services. In practice, AI-related applications are being accepted under the existing ICT regulatory sandbox.

Differences exist between the two systems. The EU sandbox focuses on pre-testing conformity with the AI Act. While participation in the

sandbox does not replace conformity assessment, receiving guidance from supervisory authorities and securing written proof containing participation results (exit report) facilitates preparation for formal market entry. Korea's regulatory sandbox is structured to temporarily exempt regulations under existing laws to test businesses.

When preparing to enter the EU market, Korean business operators may consider participating in AI regulatory sandboxes in EU member states. Countries such as France, Spain, and the Netherlands operate or are preparing to operate sandboxes. While there is no nationality restriction on participation itself, communication with the Market Surveillance Authority of the relevant member state is required.

7. Easing the Burden on SMEs and Implications for Korean Small and Medium-Sized Enterprises

Another pillar of the Omnibus revision is easing the burden on small and medium-sized enterprises (SMEs).

Although the original text of the AI Act contained provisions regarding SMEs, continuous points were raised that conformity assessment costs and documentation burdens were excessive for small and medium-sized enterprises. Industry estimates suggested that high-risk AI conformity assessments could cost tens of thousands

of euros per case, and up to hundreds of thousands of euros when combined with establishing the entire compliance framework. These are figures difficult to afford for AI startups with fewer than 50 employees.

The Omnibus amendment proposal includes directions to relax conformity assessment procedures for SMEs. While specific details will be finalized through the legislative process of the European Parliament and the Council, expanding the scope of self-declarations of conformity (self-assessments), simplified technical documentation templates, and fee reductions are being discussed.

There are two implications for Korean SMEs. Korean AI SMEs planning to enter the EU market must verify whether they are eligible for SME reduction benefits. Korean companies meeting the definition of an SME under EU law (fewer than 250 employees and an annual turnover not exceeding 50 million euros or a total balance sheet not exceeding 43 million euros) may be eligible for relief. Even without an entity established within the EU, while subject to regulation under the extraterritorial application provisions of the AI Act, they can still receive substantial portions of SME relief. However, as certain benefits—such as priority access to member state

sandboxes—require an establishment within the EU, requirements must be checked for each relief category.

Comparisons with Korean law are also meaningful. Korea's AI Basic Act does not explicitly differentiate obligations based on enterprise size regarding high-impact and high-performance AI obligations. The administrative fine ceiling of 30 million KRW effectively applies equally to both SMEs and large enterprises. Discussions regarding tiered obligations based on enterprise size may emerge in future Korean legislation as well.

8. China's AI Regulatory Timeline and Three-Country Comparison

China has not enacted a unified AI law like the EU. Instead, it sequentially created sector-specific regulations.

The Provisions on the Administration of Algorithmic Recommendations for Internet Information Services took effect in March 2022, the Provisions on the Administration of Deep Synthesis in January 2023, the Interim Measures for the Management of Generative AI Services in August 2023, and the Measures for the Identification of AI-Generated and Synthesized Content in September 2025. A comprehensive Artificial Intelligence Law remains at the stage where expert proposal drafts by academia have been made public

and mentioned as research/promotion tasks in government legislative plans, with no official government draft or enactment date released yet.

Comparing the enforcement dates of the three countries reveals an interesting pattern. China has already been applying regulations since 2022. Korea begins in January 2026. In the EU, transparency obligations apply in August 2026, and high-risk obligations apply in December 2027 (post-Omnibus).

Looking solely at the enforcement dates, Korea is 4 years behind China, but nearly 2 years ahead of the EU's high-risk regulations. However, the scope of coverage and enforcement intensity of the laws differ. China has sector-specific regulations that leave gaps, the EU is comprehensive but slow to enforce, and Korea is positioned in between.

Korean business operators operating AI services in the global market must integrate the regulatory schedules of the three jurisdictions into a single calendar. They need to be able to see at a glance which obligations apply in which country and when.

9. Documentation Unification Strategy to Prepare for Dual Regulation

A trap that business operators who must simultaneously comply with both Korean law and EU law easily fall into is dual documentation. If risk management measures to be established, maintained, and presented upon request under Korean law are prepared separately from technical documentation to be submitted to EU market surveillance authorities, management costs double.

The solution is to design the documentation structure from the outset to fulfill both sets of requirements.

Take risk management documentation as an example. Organize the requirements of Article 34(1)(1) of the Korean law (establishment and operation of risk management measures) and Annex IV of the EU AI Act (technical documentation list) into a comparison table. Map items required by Korea (risk identification, risk mitigation measures, monitoring frameworks) with items required by the EU (description of the risk management system, specifications of training, validation, and testing data, performance metrics). Overlapping items are integrated into a single section, while Korea-specific and EU-specific items are separated into distinct appendices.

It is a structure where documents for submission to Korea and for submission to the EU are extracted from a single master document.

Maintaining and managing the master document allows changes from both sides to be reflected in one place.

This strategy is suitable for medical device AI business operators. Medical device risk management reports for the Korean Ministry of Food and Drug Safety (based on ISO 14971), risk management measures under the Korean AI Basic Act, technical documentation under the EU MDR, and conformity documentation under the EU AI Act—four documents are required, but their root is a single risk management system.

The same applies to financial AI business operators. The AI guidelines of the Korean Financial Services Commission, high-impact AI obligations under the AI Basic Act, the EU's Digital Operational Resilience Act (DORA), and high-risk AI obligations under the EU AI Act—managing these four layers of regulatory documentation within a single governance framework reduces redundant work.

Documentation unification is an organizational problem, not a technical one. Legal, development, and quality management teams must work based on the same document. If each team separately creates documents tailored to regulations within its own purview, contents become inconsistent and updates get delayed. A process is

needed to designate an owner for the master document, manage change history, and periodically reflect changes in both laws.

Now that the EU's high-risk AI obligations have been postponed due to the Omnibus revision is the right time to design documentation unification. Rather than hastily creating documents after obligations enter into force, establishing the structure during the grace period is superior in terms of both cost and quality.

Chapter 26. US Colorado and Utah State Laws and NIST AI RMF: A Benchmark for Global Business Operators

In early 2026, a South Korean HR tech company with a client in Denver, Colorado, received an email: "Please send documentation confirming whether your AI job interview analysis service constitutes a high-risk AI system under the Colorado Artificial Intelligence Act (CAIA). If applicable, please also submit an impact assessment report for preventing algorithmic discrimination along with data governance documentation." At the time, the effective date of this law was scheduled for June 30, 2026. The legal team at the Seoul headquarters was unaware of the existence of this law. They knew that there was no comprehensive AI regulation law at the federal level in the United States. Their assumption that regulations would be lax in the absence of federal law was wrong. In 2025 alone, 150 AI-related laws were passed in US state legislatures. Colorado's law was enacted earlier in 2024, and employment was included among the eight core areas defined by the law. However, in May 2026, while the legal team was preparing compliance documentation, the Colorado General Assembly repealed this law and replaced it with a new one even before its effective date arrived. In the world of US state law, the rules themselves move like this.

1. The 50-State Experiment in a Country Without Federal Law

As of July 2026, the United States has not enacted comprehensive federal AI legislation. The European Union is already enforcing the AI Act, South Korea the AI Basic Act, and China the Interim Measures for the Management of Generative Artificial Intelligence Services. The world's largest AI industrial nation lacks a single corresponding federal law.

This is not to say there was no movement at the federal level. On October 30, 2023, President Biden signed Executive Order 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence." It was a document instructing federal agencies to establish AI safety guidelines, set red-teaming evaluation standards, develop deepfake detection and labeling tools, and report on AI's impact on the labor market. Because it was an executive order rather than a statute, it did not impose direct obligations on private business operators in general, except for reporting requirements requested of certain large-scale foundation model developers and computing providers. Still, it signaled that the federal government was setting the direction for AI regulation.

That signal did not last long. On January 20, 2025, President Trump revoked Executive Order 14110 on his first day in office. The Biden

administration's AI oversight framework vanished overnight. In December of the same year, President Trump signed an executive order in the opposite direction titled "Ensuring a National Policy Framework for Artificial Intelligence." It directed the Department of Justice to establish an 'AI Litigation Task Force' to legally challenge excessive state AI laws in federal courts, and ordered the Department of Commerce to identify state laws imposing undue burdens on businesses. It even included provisions linking a portion of federal budget allocations to state governments' restraint from passing AI legislation.

The federal government's attempt at preempting state laws was frustrated in Congress. The 2025 budget reconciliation bill included a provision at the House stage imposing a 10-year moratorium blocking the enforcement of state AI laws. The Senate voted 99 to 1 on July 1, 2025, to remove this provision. Both Republicans and Democrats opposed the federal government stripping state governments of their AI legislative authority.

Consequently, AI regulation in the United States is led by state legislatures. According to the Stanford HAI AI Index Report 2026, the cumulative number of AI-related laws passed in US state legislatures

from 2016 through 2025 is 62 in California, 28 in Maryland, 25 in Virginia, 24 in Utah, 18 in New York, and 17 in Texas. Only Missouri and Rhode Island have zero AI-related laws. A massive legislative experiment is underway, with the remaining 48 states regulating AI at their own paces and directions.

2. Colorado Artificial Intelligence Act: The First Comprehensive State Law Against Algorithmic Discrimination

The Colorado Artificial Intelligence Act (CAIA) is Senate Bill SB 24-205. Governor Jared Polis signed it on May 17, 2024, and its effective date was scheduled for February 1, 2026. However, this law was never enforced for even a single day.

Upon signing, Governor Polis released a letter expressing concerns. It was a warning that the compliance obligations under this law could place a burden on AI startups and small and medium-sized enterprises, thereby stifling innovation. He requested that the legislature review the law prior to its effective date to enact amendments narrowing its scope. In a special session in August 2025, the Colorado General Assembly passed a bill (SB 25B-004) postponing the effective date itself to June 30, 2026, rather than amending the law's content. Then, on May 14, 2026, before the delayed effective date arrived, Governor Polis signed SB 26-189, repealing the CAIA and

replacing it with a new law. The new law was designed around transparency and disclosure for automated decision-making technology (ADMT), with core obligations taking effect on January 1, 2027. The heavy system of obligations intended to prevent algorithmic discrimination vanished after two years without ever being enforced even once.

Even though repealed, CAIA's design remains the starting point for subsequent US state law discussions. The central concept of the law was 'consequential decision.' Eight areas were enumerated in the law: educational enrollment and opportunities, employment and hiring, financial and lending services, essential public services, healthcare services, housing, insurance, and legal services. If AI made a decision with a legal or similarly significant effect in these areas, or if AI served as a substantial factor in making that decision, the system was considered a 'high-risk AI system.'

The definition of 'algorithmic discrimination' was also present in the statutory text: unlawful or differential treatment based on protected characteristics such as age, color, disability, ethnicity, sex, or military status. In South Korea's AI Basic Act, there is no concept directly corresponding to this. While Korean law encourages autonomous

assessment of discrimination risks through fundamental rights impact assessments, it does not separately define or prohibit algorithmic discrimination.

CAIA distinguished between developers and deployers. A developer was a person doing business in Colorado that developed or intentionally modified a high-risk AI system, while a deployer was a person that deployed and used such a system. This structure corresponds to the development business operator and user business operator under Korean law.

Let us examine developer obligations first. Developers were required to provide deployers with documentation explaining the foreseeable risks and intended uses of the high-risk AI system. That documentation had to include methods for evaluating system performance, measures to prevent algorithmic discrimination, and data governance measures. Data governance measures meant the suitability of data sources, potential for bias, and bias mitigation measures. The intended output types and system monitoring methods also had to be included in the documentation. If a developer discovered that the system was reasonably likely to cause algorithmic discrimination, it was required to notify the state Attorney General

and relevant deployers within 90 days.

Deployer obligations were directed toward consumers. Deployers were required to provide prior notification to consumers that an AI system was being used. This was a basic obligation applicable even to non-high-risk AI. It was exempted only when an ordinary reasonable person would find AI involvement obvious. When high-risk AI made an adverse decision against a consumer, additional obligations applied. Deployers had to explain the reasons for the decision, the degree to which AI contributed, and the types and sources of processed data. They had to afford consumers the opportunity to correct inaccurate personal data and to appeal the adverse decision to request human review. Impact assessments had to be conducted annually, and additional impact assessments were required within 90 days if the system underwent significant modification.

Exemption provisions existed for small deployers. If a deployer employed fewer than 50 employees, did not train the AI using its own data, used the system solely for its developer-intended purpose, and made the developer's impact assessment available to consumers, it was exempt from the majority of CAIA obligations, including the impact assessment requirement. The basic obligation for prior

notification of AI usage was not exempted.

Enforcement was under the exclusive authority of the Colorado Attorney General. A private right of action allowing individuals to file lawsuits directly against companies was not recognized. The burden of proof lay on the business operator side. Developers and deployers had to affirmatively prove that they complied with CAIA requirements.

Defense mechanisms were also designed into the law. If a violation was discovered and cured through internal reviews, red-teaming tests, or external feedback, that could serve as an affirmative defense. Demonstrating compliance with the NIST AI Risk Management Framework (AI RMF), ISO/IEC 42001, or a risk management framework designated by the Attorney General could also form the basis for an affirmative defense. This provision served as the conduit that elevated the NIST AI RMF to the de facto standard of US AI regulation.

3. Utah Artificial Intelligence Policy Act: The Starting Line for Notification Obligations

In Utah, the Governor signed the Artificial Intelligence Policy Act (SB 149) in March 2024. Taking effect on May 1, 2024, it became one of the early states in the United States to enforce comprehensive AI

regulation.

Unlike Colorado, which chose the heavy obligation of prohibiting algorithmic discrimination, Utah opted for a lighter starting point. There is one core obligation: an individual or business operator interacting with consumers using generative AI must clearly disclose that the consumer is conversing with AI if requested by the consumer. Rather than requiring automatic prior notification in every interaction, it operates on a query-response model when asked by the consumer. However, under a 2025 amendment, when using generative AI in licensed regulatory occupations such as healthcare, the notification obligation was strengthened to require prior notification before the consumer asks.

Utah did not create a new regulatory agency. It enforces the law within its existing consumer protection framework. Instead, it established the Office of Artificial Intelligence Policy under the Department of Commerce to operate the 'AI Learning Laboratory' program for studying AI applications. This reflects an approach of placing understanding before regulation.

In March 2025, the Utah Legislature passed a separate bill, HB 452, regulating mental health chatbots. When AI provides emotional

counseling or mental health support services, three additional obligations apply: it must disclose to users that they are conversing with AI; it cannot sell or share users' personal health data with third parties; and if advertisements are displayed, it must clearly disclose the fact that it is an advertisement along with the sponsor, while prohibiting the use of conversation content entered by users for targeted advertising. California also enacted SB 243 in October of the same year, mandating AI identity disclosure and suicide prevention safety protocols for companion chatbot operators. This signals that mental health AI chatbots are emerging as new regulatory targets under US state laws.

Compared to South Korea's AI Basic Act, Utah's approach has a much lower regulatory intensity. Korean law mandates prior notification for products and services utilizing high-impact AI or generative AI (Article 31(1)), and imposes risk management, explanation, human intervention, and documentation obligations on high-impact AI. Utah requires disclosure only upon consumer request. Yet within the United States, Utah's law is evaluated as a pragmatic model of "let's get started," contrasting with the federal Congress, which fails to pass anything by attempting to regulate everything at once.

4. NIST AI Risk Management Framework: A Voluntary Yet Unignorable Standard

On January 26, 2023, the US National Institute of Standards and Technology (NIST) released the AI Risk Management Framework (AI RMF 1.0). It is not a statute. It is a set of voluntary guidelines without legal binding force. However, the weight this document carries in the US AI regulatory ecosystem is far heavier than what the word 'voluntary' might suggest.

The AI RMF presents the trustworthiness of AI systems through seven criteria: Validity and Reliability, Safety, Security and Resilience, Accountability and Transparency, Explainability and Interpretability, Privacy, and Fairness. These seven criteria overlap significantly with the fundamental principles of South Korea's AI Basic Act (Article 3) and the requirements of the EU AI Act. This means an international common language for AI regulation has already been formed.

The execution structure of the framework consists of four core functions.

Govern encompasses the whole framework. It is the function of establishing and maintaining a governance structure for AI risk management across the entire organization. It addresses policies, procedures, allocation of responsibility, and organizational culture.

For the remaining three functions to operate in the field, governance must be established first. South Korea's AI Basic Act requiring business operators to establish and operate risk management measures (Article 34(1)(1)) and to publish the name and contact information of the person managing and supervising high-impact AI (Enforcement Decree Article 27(1)) corresponds to this function.

Map contextualizes risks. It analyzes the environment and conditions in which the AI system will be deployed to identify and categorize risks. It asks who is impacted by the system, under what conditions the system could fail, and where bias might be introduced. The procedure under Korean law for determining high-impact AI—namely, the structure for verifying domain applicability as well as the impact and severity of risk (Act Article 33, Enforcement Decree Article 25)—resembles this function.

Measure analyzes and tracks identified risks through quantitative metrics and qualitative assessments. It continuously monitors the magnitude, frequency of risks, and changes in system performance. Korean law requiring high-impact AI business operators to periodically improve risk management measures corresponds to this function.

Manage allocates resources and responds to risks. Based on the results of mapping and measuring, it plans and executes concrete actions to decrease risk. It determines whether to accept, transfer, mitigate, or avoid risks.

The four functions are not a one-time checklist, but an iterative cycle applied continuously throughout the lifecycle of the AI system.

There was a decisive turning point that prevented the AI RMF from remaining merely voluntary guidance: the Colorado Artificial Intelligence Act. CAIA was designed to grant affirmative defense eligibility to business operators complying with the NIST AI RMF. It served as a legal defense mechanism enabling a business operator facing alleged violations to state, "We operated our risk management system in accordance with the NIST AI RMF." Although the law itself was repealed prior to enforcement, it remains a landmark example demonstrating how a voluntary framework can be converted into a de facto legal standard through state law.

In 2024, NIST released an additional Generative AI Profile (AI RMF Generative AI Profile), supplementing management methods for risks unique to large language models and multimodal AI. Hallucination, copyright infringement, harmful content generation, CBRN (chemical,

biological, radiological, nuclear) risks, and information security threats fall under its scope. If the main AI RMF text is the framework for general AI risk management, the Generative AI Profile weaves thread specific to generative AI onto that framework.

5. Coexistence of Federal Soft Law and State Hard Law

A unique dual structure exists in the US AI regulatory landscape.

At the federal level, the NIST AI RMF, executive orders, and agency-specific guidance exist as soft law. They lack legal binding force or, like executive orders, can be revoked when the presidency changes. The case of the Biden administration's AI executive order vanishing on day one of the new administration is evidence of this. As long as Congress fails to pass comprehensive AI legislation, AI norms at the federal level remain inherently unstable.

At the state level, statutes operate as hard law. The final trajectory of Colorado's CAIA demonstrates this character. Even after the federal Department of Justice joined unconstitutionality litigation and a federal court issued a preliminary injunction, what actually repealed the law in the end was the Colorado General Assembly itself. The amendment and repeal of state laws rest, in principle, in the hands of state legislatures. The Illinois Biometric Information Privacy Act

(BIPA) has remained effective for 18 years since taking effect in 2008. State laws persist independently of presidential transitions.

The nexus where these two tiers met was CAIA's affirmative defense provision. It was a structure under which complying with the federal voluntary framework (NIST AI RMF) provided a legal defense against violations of state law. Soft law was incorporated as a compliance standard for hard law. As more state laws reference the NIST AI RMF going forward, this voluntary framework draws closer to becoming a de facto national standard. It becomes a standard that is difficult to ignore even though it is not a statute.

The Trump administration's December 2025 executive order added tension to this structure. The Department of Justice's AI Litigation Task Force declared it would challenge "unduly burdensome" state laws in federal courts. However, as demonstrated by the Senate's 99 to 1 vote, both parties oppose the federal government stripping state governments of their AI legislative power. While child safety in states, data center infrastructure, and state procurement are cited as areas retaining room for state regulation, boundaries for remaining areas remain unclear. This tension persists unresolved.

The regulatory frameworks of the EU and South Korea are fundamentally different from this dual structure. The EU AI Act is a Regulation directly applicable across all 27 member states. Room for member states to enact conflicting laws is limited. South Korea's AI Basic Act applies uniformly nationwide. The rules are identical whether in Seoul or Jeju. Only the United States operates a disparate AI regulatory system state by state.

6. South Korean Business Operators Entering the US: Reading State-by-State Regulatory Maps

When South Korean AI companies enter the US market, the regulatory environment they face varies state by state. The same product encounters different rules each time it crosses a state border.

As of July 2026, state-by-state circumstances fall into three groups: enacted and effective, pending enforcement, and repealed/replaced. To operate an AI chatbot service in Utah, one must immediately disclose the AI identity upon consumer request; if it is a mental health chatbot, selling or sharing health data with third parties and targeted advertising based on user input are prohibited. Some states are awaiting enforcement. California's AI Transparency Act (SB 942) requires generative AI providers above a certain scale to provide watermarking and free detection tools starting August 2, 2026, while

Washington State's HB 1170 requires providers of generative AI systems with over 1 million monthly active users to provide provenance data for AI-generated media starting February 1, 2027. Then there are states where laws have been repealed and replaced. Colorado's CAIA, which intended to mandate annual impact assessments to prevent algorithmic discrimination for AI hiring tools, vanished before enforcement, and a new law centered on transparency and disclosure will apply starting January 1, 2027.

Illinois requires separate attention. The Biometric Information Privacy Act (BIPA), in effect since 2008, is not an AI-specific law, but applies directly to facial recognition AI. If biometric information is collected without an individual's prior written consent, individuals can file lawsuits directly against companies. Meta settled for \$650 million (approx. 850 billion KRW) over allegations of collecting facial recognition data without consent in its Facebook photo tagging feature. Clearview AI received judicial approval in 2025 for a settlement forfeiting 23% of company equity (valued at approximately \$51.75 million); however, in July 2026, an appellate court revoked approval due to failure to meet class representative requirements, remanding the case to the court of first instance. Unlike Colorado's CAIA, which placed enforcement exclusively under

the Attorney General, Illinois allows individuals to launch class action lawsuits. The dimension of litigation risk is entirely different.

Texas shows a different aspect. The Responsible AI Governance Act (HB 149) was passed in 2025 and took effect in January 2026, but significantly reduced private sector obligations compared to the original bill, focusing on prohibiting behavioral manipulation, creation of child sexual abuse material, intentional discrimination, and infringement of constitutional rights. Under the common title of 'AI regulation,' the breadth and depth of regulation vary completely from state to state.

The contrast with South Korea's regulatory environment is stark. In South Korea, a single law, the AI Basic Act, applies nationwide. The rules are the same from Seoul to Busan. To do business in the United States, one must individually verify the laws of every state reached by the service.

A realistic response follows two paths. One is using the standard of the state with the highest regulatory baseline as a common foundation. However, because each state law differs regarding scope of application, timing of notification, and rules on health data, biometric data, and advertising, a single strict standard will not

automatically resolve requirements across all states, necessitating separate verification of state-specific rules. The other path is mapping out state-by-state regulatory requirements and tailoring services to each state's mandates. While costlier, it enables a wider operational scope.

Under both paths, establishing a risk management framework based on the NIST AI RMF serves as the foundation. This is because the repealed Colorado law established a precedent of incorporating this framework as an affirmative defense standard, and other state laws are likely to make similar references in the future.

7. South Korea, the EU, and the US: Comparison of Three Regulatory Philosophies

The three systems depart from different questions.

The EU AI Act asks: "Does this AI system pose a risk of infringing upon fundamental rights?" If the risk is unacceptable, it is prohibited. If high-risk, it must undergo a prior conformity assessment before entering the market. It is a single Regulation directly applicable across all 27 member states, with violations of the most severe prohibited practices subject to administrative fines of up to €35 million or 7% of total worldwide annual turnover. Social scoring and

subliminal manipulation are banned, and real-time remote biometric identification in publicly accessible spaces for law enforcement purposes is also prohibited except under narrow exceptions such as searching for missing persons. Regulation is clear and strong. Consequently, prior burdens on business operators are heavy.

South Korea's AI Basic Act asks: "Does this AI system exert a significant impact on society?" It imposes risk management and transparency obligations on high-impact AI, but contains no prior conformity assessments or categories of prohibited practices. It employs soft regulation such as autonomous effort duties and public procurement incentives. Its conditions include uniform nationwide application, a maximum administrative fine of 30 million KRW, and a grace period of over one year announced by the government. It is a hybrid approach embedding innovation promotion and rights protection within a single statute.

The United States has no unified question at the federal level. Colorado asked, "Does this AI discriminate?" before changing the question to "Does the consumer know about the automated decision?" upon replacing the law; Utah asks, "Does the consumer know it is AI?"; and Illinois asks, "Was biometric consent obtained?"

The NIST AI RMF asks, "Has a risk management system been established?", though it lacks legal binding force. The 50 states are crafting their own questions and answers.

Let us quantify enforcement power in figures. The EU's maximum fine is €35 million, approximately 59.5 billion KRW at 2026 exchange rates. China can impose up to 50 million RMB (approx. 10 billion KRW) or 5% of the preceding year's turnover under its Personal Information Protection Law alone, whichever is higher. South Korea caps administrative fines at 30 million KRW. The United States varies by state: Colorado focuses on Attorney General sanctions, while Illinois BIPA can generate multi-hundred-million-dollar settlements through class actions via private rights of action, making practical risk far exceed statutory caps.

Something is clear from this comparison: in terms of systematic protection of fundamental rights, the EU leads among the three. In balancing industrial growth velocity against business operator burdens, South Korea's soft regulation holds an advantage. In terms of experimentation and diversity, the US state-by-state decentralized model functions as a massive laboratory testing 50 different approaches simultaneously. One cannot conclusively determine

which of the three models is correct. However, for South Korean business operators, the only certainty is the reality that they must read and prepare for all three regulatory environments.

If a South Korean AI business operator enters both the US and the EU simultaneously, it must operate a triple-compliance system: establishing technical documentation and a risk management system at the level of the EU's prior conformity assessment; meeting state-specific additional requirements in the US based on the NIST AI RMF; and executing transparency obligations and high-impact AI determination procedures under South Korea's AI Basic Act. Because the EU regulatory standard is the highest among the three, adopting EU compliance as the baseline and supplementing gaps for South Korea and the US is the realistic approach.

Although this triple structure poses a burden, it is not necessarily disadvantageous for South Korean business operators. This is because South Korea's AI Basic Act shares many core principles with the EU AI Act and the NIST AI RMF. The four pillars of transparency, risk management, human oversight, and fairness are common to all three systems. A business operator that has faithfully complied with Korean law already possesses a solid foundational starting point for

adapting to US and EU requirements. The remaining task is to accurately measure the distance from that starting base to the destinations of each market and bridge those gaps.

Part 9 Practices for Business Operators by Industry

Chapter 27. Financial AI Business Operators

In the autumn of 2023, a loan application by a self-employed business owner in their 30s was rejected at a mutual savings bank in Yeouido, Seoul. The counter employee could not explain the reason. They could only show a screen indicating that the credit score fell below the threshold. It was not a human, but an AI credit scoring model that assigned that score. Neither the employee nor the bank's headquarters could provide a clear answer regarding what data was entered and by what criteria the score was deducted. Although this self-employed business owner had the right to request an explanation of automated evaluation results under the 2020 amendment to the Credit Information Act, they were not even informed that such a right existed.

Financial AI differs in character from AI in other industries. A machine determines or recommends whether one can obtain a loan, subscribe to insurance, or where to invest funds. If something goes wrong, a person's entire economic life is shaken. Therefore, financial AI business operators are subject not only to the AI Basic Act but simultaneously to separate guidelines from the Financial Services Commission and regulations under the Credit Information Act. Let us

unpack step-by-step how this structure, often called dual regulation, imposes burdens on financial AI business operators and what protective mechanisms it establishes.

1. Structure of Dual Regulation: Financial Services Commission Guidelines and the AI Basic Act

AI regulation in the financial sector consists of two layers. On one side are the "Financial Sector AI Guidelines," first released by the Financial Services Commission in July 2021 and revised in June 2026; on the other side is the AI Basic Act, which entered into force on January 22, 2026.

The Financial Services Commission Guidelines are not an Act of law. They are administrative guidance recommended to financial companies and fintech firms by an administrative agency. Although they lack legal binding force, compliance with these guidelines functions as a de facto evaluation item during Financial Supervisory Service examinations. From the perspective of financial companies, they carry practical binding power that makes it impossible to ignore the recommendations.

The AI Basic Act is a law. Whether in finance, healthcare, or education, it applies to any AI business operator. Article 5(1)

stipulates that "Except as otherwise specially provided in other Acts, the provisions of this Act shall apply." If AI-related provisions already exist in the Credit Information Act, those provisions take precedence; if not, the AI Basic Act applies. Because the Financial Services Commission Guidelines are not an Act of law, they do not possess the status of a special law over the AI Basic Act.

Operationalizing this structure in practice yields the following: a bank operating an AI loan screening system must explain the main criteria and overview of basic information of automated evaluations to users pursuant to Article 36-2 of the Credit Information Act. Simultaneously, under Article 2, Item 4(g) of the AI Basic Act, this system may constitute high-impact AI, and if so, additional obligations under Article 34(1) arise: establishing risk management measures, explaining grounds for decisions, preparing user protection measures, securing human intervention systems, and retaining documentation for five years. On top of this, the Financial Services Commission Guidelines demand governance frameworks, the principle of auxiliary role, and the securing of reliability.

2. Key Principles of the Financial Services Commission AI Guidelines

When first published in July 2021, the Financial Services Commission Guidelines set forth five principles that financial

companies must observe when adopting and operating AI. The revised version, announced on June 18, 2026, and effective from June 22, reorganized these into seven principles: Governance, Legality, Auxiliary Role, Reliability, Financial Stability, Good Faith, and Security. The first three principles have been maintained from the beginning, while the previous principles of explainability and fairness were removed from the titles of independent principles, though their requirements are integrated into the new framework of principles.

First is Governance. This requires establishing internal control systems for the development and operation of AI systems. The AI utilization strategy must be approved at the board of directors or executive management level, and a designated officer responsible for managing AI-related risks must be named. In effect, AI risk oversight is added to the roles of the Chief Information Security Officer (CISO) or Chief Risk Officer (CRO). This poses a burden on small and medium-sized fintech companies. How startups lacking the capacity to employ dedicated personnel will fulfill the demand to establish board-level governance remains a practical challenge.

Second is Legality. This principle requires compliance with relevant laws and regulations when utilizing AI. While it sounds self-evident,

the scope of applicable laws in the financial sector is broad. The Credit Information Act, Electronic Financial Transactions Act, Capital Markets Act, Insurance Business Act, and Personal Information Protection Act overlap. Operating a single AI model requires auditing the requirements of five separate laws simultaneously.

Third is Auxiliary Role. This principle dictates that an AI's decision must serve as an auxiliary tool to assist human judgment rather than a final decision. Even if AI rejects a loan approval, a structure must be maintained in which a human reviews the result and makes the final determination. This aligns with the human management and supervision required by Article 34(1)(4) of the AI Basic Act. In practice, the issue lies in superficial human intervention. When AI outputs a "rejection," employees rarely overturn the judgment in practice, beyond refraining from pressing an approval button. For the principle of auxiliary role to function effectively, authority and standards enabling employees to overturn AI results must be established concurrently.

Fourth is Reliability. This is the principle that AI decision outcomes must be trustworthy and usable, with explainability at its core. It must be possible to explain to users why AI produced a particular

result. Article 36-2 of the Credit Information Act is a provision that codifies this requirement into law. Individuals may demand explanations of automated evaluation results, primary criteria, and overviews of basic information. Where basic information is inaccurate or outdated, individuals may request correction or deletion, or request recalculation of results. Going a step further than the law, the Financial Services Commission Guidelines recommend efforts to explain the decision logic of the AI model itself—recommending the application of post-hoc explainability techniques even when black-box models are used.

Requirements for fairness audits also continue. It must be audited to ensure that AI does not produce discriminatory results based on factors such as gender, age, race, or region of origin. Credit scoring models systematically assigning lower scores to residents of specific regions, or premium calculation models introducing unjustifiable discrepancies based on gender, represent typical violations of fairness. If historical discrimination patterns are embedded in training data, AI reproduces those patterns. The guidelines mandate auditing bias in training data and monitoring the fairness of outcomes.

The principles newly established in the revised version are Financial Stability, Good Faith, and Security. These reflect awareness that AI malfunctions can destabilize the financial market as a whole beyond individual customer damages, reaffirmation that financial companies retain duty of good faith toward customers even when utilizing AI, and demands to prepare against security threats targeted at AI systems.

3. AI Credit Scoring and Loan Screening

AI credit scoring is a use case with a long history in financial AI. Credit bureaus (CBs) such as NICE Information Service and Korea Credit Bureau (KCB) operate AI models that calculate individual credit scores by inputting hundreds of variables. Banks, credit card companies, and mutual savings banks utilize these scores in loan screening.

The Credit Information Act directly regulates this realm. Article 36-2 stipulates rights to explanation and objection regarding automated evaluations. "Automated evaluation" refers to evaluation processed solely by information processing devices such as computers without employees of credit information companies or similar entities participating in evaluation tasks. Pursuant to this clause, an individual credit information subject may request four things:

whether automated evaluation is being conducted, the results of the automated evaluation, the main criteria, and an overview of basic information.

Added to this is the AI Basic Act. Loan screening is explicitly listed in Article 2, Item 4(g): "Judgments or evaluations that significantly affect an individual's rights and obligations, such as employment and loan screening." An AI loan screening system constitutes high-impact AI. High-impact AI business operators are imposed with obligations under Article 34(1): establishing and operating risk management measures (Item 1), establishing and implementing explanation measures (Item 2), establishing and operating user protection measures (Item 3), human management and supervision (Item 4), and drafting and retaining documentation evidencing performance (Item 5). Detailed matters such as posting key contents on websites and retaining documentation for five years are prescribed by Article 27 of the Enforcement Decree.

The position of credit bureaus such as NICE Information Service is complex. NICE develops AI credit scoring models and supplies them to financial companies, making it a development business operator under the AI Basic Act. Simultaneously, NICE itself directly bears

explanation obligations for automated evaluations as an individual credit bureau under the Credit Information Act. Financial companies (banks, credit card companies) receive NICE's models and utilize them for loan screening, making them deployer business operators under the AI Basic Act. If NICE, as the development business operator, has fulfilled the risk management measures, explanation measures, and user protection measures, those three measures are deemed to have been fulfilled by the bank as the deployer business operator unless the bank substantially modifies the model's functionality (Article 27(3) of the Enforcement Decree). However, human management and supervision, as well as document preparation and retention, must be separately fulfilled by the bank. This is a mechanism to prevent duplicate burdens.

Prevention of duplicate regulation is also stipulated in the Enforcement Decree of the AI Basic Act. Annex 1 of the Enforcement Decree specifies item by item which duties under other laws, if complied with, are deemed to fulfill which obligations under the AI Basic Act. If an entity complies with explanation duties under Article 35-2 of the Credit Information Act and maintains explanation procedures pursuant to Article 36-2(1), the obligation for explanation measures under Article 34(1)(2) of the AI Basic Act is recognized as

fulfilled. However, obligations not included in Annex 1—such as establishing risk management measures or retaining documentation for five years—must be separately fulfilled in accordance with the AI Basic Act.

4. Robo-Advisors: When AI Recommends Investments

Robo-advisors are services in which algorithms analyze an investor's inclination and goals to construct portfolios and automatically rebalance them. In South Korea, the Financial Services Commission has operated a robo-advisor testbed since 2016. For certain business modes, such as non-face-to-face discretionary investment management, passing the testbed operated by Koscom functions as a de facto prerequisite.

The Capital Markets Act regulates investment advisory and discretionary investment management businesses. If a robo-advisor "recommends" investments, it falls under investment advisory; if it directly manages funds entrusted by customers, it falls under discretionary investment management. In the case of discretionary investment management, operations must align with each customer's risk profile, and the potential for loss must be fully disclosed. Managing via AI does not exempt an entity from these obligations.

The point of intersection with the AI Basic Act is as follows: The key issue is whether the act of a robo-advisor determining or recommending an individual's asset allocation falls under "judgments that significantly affect an individual's rights and obligations" under Article 2, Item 4(g). Unlike loan screening, it is not explicitly listed in the statutory text. Room for interpretation exists regarding whether it is encompassed by the term "such as." Erroneous investment recommendations result in major financial impacts, namely loss of principal. There is a high probability it will be classified as high-impact AI. However, as of July 2026, no official authoritative interpretation from the Ministry of Science and ICT has been confirmed.

When applying the principle of auxiliary role under the Financial Services Commission Guidelines to robo-advisors, a process in which a human confirms the final trading decision is required even if AI constructs the portfolio. Tension exists between fully automated discretionary investment management (where algorithms automatically execute buys and sells) and the principle of auxiliary role. The Financial Services Commission has not explicitly resolved this tension.

5. AI Insurance Screening

In the insurance industry, AI is utilized in two primary areas: underwriting and claims processing.

In underwriting, AI analyzes insurance applicants' health information, occupations, lifestyle habits, and prior claim histories to determine eligibility and premium rates. AI replaces a significant portion of the tasks traditionally performed by actuaries. In claims processing, it determines the authenticity of submitted claim documents and scores the likelihood of fraud.

Although the Insurance Business Act regulates authorization of insurance products and review of terms and conditions, it contains no provisions directly regulating AI underwriting itself. The Financial Services Commission Guidelines fill this void. Under the principle of fairness, audits must be conducted to ensure AI underwriting models do not unjustly overweight specific diseases or genetic information.

Regarding the relationship with the AI Basic Act, interpretations differ on whether insurance underwriting constitutes high-impact AI. Article 2, Item 4(g) states "judgments or evaluations that significantly affect an individual's rights and obligations, such as employment and loan screening." Rejection of insurance coverage significantly affects an individual's economic rights; viewing it as included under "such

as" is a natural interpretation. The same applies to claims processing. A decision denying insurance payouts directly impacts an individual's property rights.

The reason AI fairness issues become contentious in insurance stems from the characteristics of training data. Insurance is inherently a business of categorizing risk. Charging higher premiums to high-risk groups is the principle of insurance. However, AI can execute de facto discrimination under the rubric of "risk classification." A model systematically assessing higher premiums on residents of specific postal codes becomes a tool for regional discrimination. The boundary between risk classification and prohibition of discrimination in insurance was controversial even prior to AI, but AI has rendered that boundary even more blurred.

6. Financial AI as High-Impact AI: Article 2, Item 4(g)

Article 2, Item 4 of the AI Basic Act enumerates 11 application areas for high-impact AI. Item (g) is "judgments or evaluations that significantly affect an individual's rights and obligations, such as employment and loan screening."

Applying this clause to financial AI, loan screening is explicitly included. Employment is also included. The phrase "such as"

expands the scope. Credit scoring, insurance underwriting, claims settlement evaluations, and investment suitability assessments may fall under "judgments that significantly affect an individual's rights and obligations."

When classified as high-impact AI, five duties are imposed on business operators under Article 34(1).

Business operators must establish and operate risk management measures. A cyclical framework consisting of policy formulation, risk identification, analysis and evaluation, treatment, and improvement must be established. Because financial companies already maintain operational risk management frameworks, incorporating AI risks into existing frameworks is a realistic approach.

Business operators must establish and implement explanation measures. Within technically feasible limits, they must be capable of explaining the final results derived by AI, the main criteria used in judgment, and an overview of training data. Pursuant to Article 27(1) of the Enforcement Decree, the main contents of risk management measures, explanation measures, user protection measures, and the name and contact details of the supervisor/manager must be posted at the workplace or on the website, with trade secrets eligible for

exclusion.

Business operators must establish and operate user protection measures. This requires setting up mechanisms to prevent and respond to harm that users might suffer while using AI products or services.

Business operators must establish systems allowing human intervention. Procedures are required whereby humans review AI decisions and overturn them if necessary. This shares the same orientation as the principle of auxiliary role under the Financial Services Commission Guidelines.

Business operators must retain documentation evidencing performance for five years (Article 27(2) of the Enforcement Decree). This period overlaps with financial companies' general record-keeping obligations (typically 5 to 10 years).

7. Dual Regulation Burden and Mechanisms to Prevent Duplication

The burden felt by financial AI business operators is as follows: In operating a single AI loan screening system, explanation duties for automated evaluation under the Credit Information Act, high-impact AI duties under the AI Basic Act, the seven principles of the Financial

Services Commission Guidelines, and automated decision-making provisions under the Personal Information Protection Act overlap. The core practical question is which clause of which law, if fulfilled first, can be deemed to satisfy obligations under another law.

The AI Basic Act incorporates mechanisms to prevent duplicate regulation. Article 5(1) stipulates that "Except as otherwise specially provided in other Acts, the provisions of this Act shall apply." Where AI-related provisions exist in the Credit Information Act, the Credit Information Act takes precedence.

The Enforcement Decree follows the same direction. Article 27(5) and Annex 1 of the Enforcement Decree specify item by item which duties under existing legislation—such as the Personal Information Protection Act, Digital Medical Products Act, and Credit Information Act—if complied with, are deemed to fulfill which obligations under the AI Basic Act.

Summarizing which specific duties are mutually recognized yields the following:

Explanation duties are mutually recognized. If an entity complies with explanation duties under Article 35-2 of the Credit Information

Act and maintains explanation procedures under Article 36-2(1), it need not separately fulfill the obligation for explanation measures under Article 34(1)(2) of the AI Basic Act (Annex 1, Item 2(c)).

User protection measures are also mutually recognized. If an entity fulfills all duties under Article 10 of the Financial Consumer Protection Act, it is deemed to have fulfilled user protection measures under Article 34(1)(3) of the AI Basic Act (Annex 1, Item 3(b)).

Rights to object and request correction present a different situation. Article 36-2(2) of the Credit Information Act stipulates rights to request correction or deletion of basic information and recalculation, overlapping in purpose with human management and supervision requirements under the AI Basic Act. However, because Annex 1 does not specify this as a deemed-performance item, duties are not exempted merely because their underlying purposes overlap.

Areas without mutual recognition also exist. Annex 1 exhaustively enumerates recognized items, but contains no item recognizing risk management plan establishment, human management and supervision, or five-year document retention with respect to loan screening. These obligations, including website posting, must be fulfilled separately.

Because the Financial Services Commission Guidelines are not an Act of law, they are not subject to mutual recognition with the AI Basic Act. No matter how faithfully guidelines are fulfilled, obligations under the AI Basic Act are not exempted. Conversely, even if one complies with the AI Basic Act, ignoring recommendations in the guidelines can draw criticism during Financial Supervisory Service examinations.

8. Article 36-2 of the Credit Information Act: Explanation and Objections Regarding Automated Evaluation

Article 36-2 of the Credit Information Act constitutes the substantive core of financial AI regulation. Newly established during the 2020 Data 3 Acts revision, it has been in effect since August 5, 2020.

"Automated evaluation," the subject of this clause's application, falls into two categories: individual credit evaluation (credit score calculation) and decisions on the establishment, maintenance, or contents of financial transactions (loan approval, limit setting, etc.). Applicable business operators include individual credit bureaus (CBs) and financial companies prescribed by Presidential Decree. Pursuant to Article 31-2 of the Enforcement Decree, this encompasses entities specified under each item of Article 38 of the Act on the Establishment, etc. of Financial Services Commission—meaning all

financial institutions subject to examination by the Financial Supervisory Service, including banks, insurers, and securities firms.

The scope of explanation requests is broad. Individuals may request explanations on four items: whether automated evaluation is conducted, results, main criteria, and an overview of basic information.

Rights to object also exist. Individuals may submit information favorable to themselves, and if basic information is inaccurate, request correction or deletion, or request recalculation of results. The right to request recalculation effectively means being able to obtain a re-examination of AI decisions.

Grounds for refusal are limited to three cases: where special provisions exist in this Act or other Acts, or where unavoidable to comply with statutory duties; where establishing and maintaining commercial relationships such as financial transactions becomes difficult; or similar cases prescribed by Presidential Decree. Because grounds for refusal are exhaustively enumerated, explanation cannot be refused unless one of these three cases applies.

Compared to explanation duties under the AI Basic Act (Article 34(1)(2)), the Credit Information Act is more specific and stringent. Article 34(1)(2) of the AI Basic Act stipulates preparing explanation measures "within technically feasible limits," granting business operators broad discretion. Article 36-2 of the Credit Information Act specifies items requiring explanation (results, main criteria, basic information overview) and exhaustively enumerates grounds for refusal. For financial AI business operators, the requirements under the Credit Information Act serve as the practical benchmark.

9. Position of Financial AI in the EU AI Act

The European Union AI Act explicitly classifies financial AI as high-risk.

Annex III, Paragraph 5 covers the area of "Access to and enjoyment of essential private services and essential public services and benefits." This includes creditworthiness evaluation and credit scoring of natural persons. Risk assessment and pricing in relation to natural persons in the case of life and health insurance are also classified as high-risk AI.

Korea focuses on impact; the European Union focuses on risk. This difference yields concrete outcomes in financial AI. To launch an AI

credit scoring system in the EU, one must pass a prior conformity assessment. One must draft technical documentation, establish a risk management system, implement data governance, and set up human oversight measures. Registration in the EU database is required. However, procedures to defer the application date of these high-risk obligations from the original August 2026 to December 2027 are effectively in their final stages. Violations carry fines of up to 15 million euros or 3% of global turnover, whichever is higher.

In Korea, even if the same AI credit scoring system is classified as high-impact AI, there is no requirement to undergo prior conformity assessment. Business operators need only voluntarily establish and operate risk management measures. Violations do not immediately trigger administrative fines; if an entity fails to comply even after receiving a corrective order from the Ministry of Science and ICT, an administrative fine of up to 30 million KRW is imposed (Article 40(3), Article 43(1)). The disparity in regulatory intensity between the EU and Korea is starkly evident in financial AI.

The EU AI Act is also specific in the insurance field. Using AI for risk assessment in life and health insurance is classified as high-risk. Auto insurance and property insurance are excluded. Only insurance

directly related to human life and health is deemed high-risk. South Korea's AI Basic Act makes no distinction by insurance type. Under the comprehensive criterion of "judgments that significantly affect an individual's rights and obligations," systems may fall under high-impact AI regardless of insurance type.

10. Compliance Practice for Financial AI Startups

The reality faced by fintech firms in Yeouido and Gangnam is as follows:

Regulations are numerous. The AI Basic Act, Credit Information Act, Personal Information Protection Act, Capital Markets Act (investment-related), Electronic Financial Transactions Act (payment-related), and Insurance Business Act (insurance-related). Depending on the business model, three or four laws apply simultaneously. On top of that, the Financial Services Commission Guidelines are layered.

There are also advantageous aspects for startups. The grace period (guidance period) of the AI Basic Act is announced to be at least one year and can be extended if necessary. Although legal obligations themselves arise during the grace period, administrative guidance and corrective recommendations take precedence over sanctions

such as administrative fines for omissions in transparency duties or insufficient documentation. Anonymous consulting is available at the AI Basic Act Support Desk. For smaller startups, it is practical to leverage this grace period to gradually build compliance systems.

The division of roles between development business operators and deployer business operators also reduces burdens. If a fintech firm developing an AI model fulfills risk management measures, explanation measures, and user protection measures as a development business operator, the financial company adopting and using that model (deployer business operator) does not need to fulfill these three measures separately unless it substantially alters the model's functionality. Human management and supervision as well as document preparation and retention must be separately fulfilled by deployer business operators themselves. From the perspective of fintech firms, completing compliance as a development business operator serves as a sales competitive edge toward financial companies.

The most difficult point is that interpretations of the law and guidelines have not yet stabilized. Authoritative interpretations from the Ministry of Science and ICT have not been issued regarding

whether "judgments that significantly affect an individual's rights and obligations" encompass asset management recommendations, premium calculations, or cross-border remittance risk evaluations. While authority for fact-finding investigations and corrective orders under the AI Basic Act resides with the Ministry of Science and ICT (Article 40), how roles will be divided and coordinated with the Financial Services Commission and Financial Supervisory Service that supervise financial companies remains unsettled.

11. Blind Spots in Financial AI Regulation

Two blind spots are visible.

Article 36-2 of the Credit Information Act targets "individual credit evaluation" and "decisions on the establishment, maintenance, or contents of financial transactions." Insurance underwriting and claims processing are not explicitly listed. The Insurance Business Act also contains no direct provisions regarding AI screening. AI in the insurance sector sits in a legal void, with only the Financial Services Commission Guidelines loosely filling that vacuum.

Regulation regarding cases where AI is used in high-frequency trading (HFT) and algorithmic trading is also insufficient. While the Capital Markets Act contains some provisions on algorithmic trading,

comprehensive regulation regarding systemic risks AI poses to the market (events such as flash crashes) is absent. The EU AI Act likewise did not classify this area as high-risk, because financial market stability belongs to a dimension distinct from infringement on individual rights and obligations. However, as demonstrated by the 2010 US Flash Crash and the 2012 Knight Capital incident, malfunctions in AI algorithmic trading can shake the entire market.

12. What Financial AI Business Operators Should Do Now

The checklist for business operators running financial AI is as follows:

Verify whether an automated evaluation explanation system under Article 36-2 of the Credit Information Act is in place. Procedures and systems must exist to provide results, main criteria, and basic information overviews upon user request. This is the starting point for financial AI compliance.

Determine whether one's AI system constitutes high-impact AI under Article 2, Item 4(g) of the AI Basic Act. If used for loan screening, it applies. If used for credit scoring, there is a high probability it applies. If used for insurance underwriting or investment suitability assessment, there is a high probability it applies.

If classified as high-impact AI, establish risk management measures. Adding AI risk items to existing operational risk management frameworks is an efficient approach. Build a framework to retain documentation evidencing performance for five years.

Audit the seven principles of the Financial Services Commission Guidelines: Governance, Legality, Auxiliary Role, Reliability, Financial Stability, Good Faith, and Security. Evaluate the firm's current level of performance against each principle and remedy deficient areas.

Secure communication channels with regulatory authorities before the law takes firm root. Anonymous consulting at the AI Basic Act Support Desk is free of charge. The Financial Supervisory Service Fintech Center also accepts inquiries. In areas with uncertain interpretation, asking proactively is better than receiving sanctions later.

Financial AI regulation moves according to two clocks. The clock of the Credit Information Act has already been running for six years. The clock of the AI Basic Act started on January 22, 2026. Where both clocks point in the same direction, a single implementation can satisfy both laws. Where directions differ, each must be addressed separately. Discerning which clock points in which direction is the

first step of compliance for financial AI business operators.

☐☐☐ ☐☐☐☐ ☐☐☐ ☐☐☐☐ ☐☐, ☐☐☐ ☐☐ ☐☐. ☐☐ ☐☐☐☐
☐☐☐ ☐☐ ☐☐ • ☐☐ • ☐☐☐☐

Chapter 28. Healthcare and Bio AI Business Operators

It is the department of radiology at a university hospital in Seoul. On a Monday morning in February 2026, a radiologist opened a chest CT scan. On the right side of the screen, AI reading-assistance software is marking three areas suspected of lung nodules with yellow circles. This software received Class 3 medical device approval from the Ministry of Food and Drug Safety (hereinafter MFDS). At the same time, according to the AI Basic Act that came into effect half a month ago, AI systems used in the healthcare provision and usage system may fall under high-impact artificial intelligence. A practical question comes to the specialist's mind: Having obtained MFDS approval, do I still need to separately fulfill the obligations under the AI Basic Act? Am I a business operator or a user under this law?

1. Structure of Dual Regulation

Healthcare AI business operators are placed within two regulatory systems.

One is the MFDS medical device approval review system. Medical device laws and regulations divide medical devices into Class 1 through Class 4 depending on risk. Software as a Medical Device (SaMD) equipped with AI is mostly classified into Class 2 or Class 3.

Products classified as Class 3 or higher can only be used on patients after obtaining approval through the MFDS's technical document review and clinical trial data examination. Approved AI medical devices by the MFDS totaled 64 cases in 2023 alone, increasing to 108 cases in 2024.

The other is the regulation of high-impact AI under the AI Basic Act. Article 2, Item 4 of the AI Basic Act enumerates 11 application domains of high-impact AI from Subitem (a) to Subitem (k), with the medical field appearing in two places. Subitem (c) covers the establishment and operation of healthcare provision and usage systems pursuant to the Framework Act on Health and Medical Services. Subitem (d) covers the development and use of medical devices pursuant to the Medical Devices Act and digital medical devices pursuant to the Digital Medical Products Act.

The two regulations overlap. AI image-reading software approved by the MFDS is regulated by the Medical Devices Act while simultaneously being classified as high-impact AI under the AI Basic Act. Development business operators must submit technical documents and clinical data to the MFDS, while also fulfilling obligations regarding risk management, explanation, user protection,

human management and supervision, and documentation as specified in Article 34, Paragraph 1 of the AI Basic Act.

2. The MFDS Approval System: How AI Medical Devices Enter the Market

The MFDS review of AI medical devices has points that differ from traditional medical device reviews.

Traditional medical devices retain their approval once approved as long as the design does not change. For AI medical devices, performance changes when training data is added or algorithms are updated. To respond to this issue, the MFDS has revised its approval review guidelines multiple times. On January 24, 2025, it published the world's first approval review guidelines for generative artificial intelligence medical devices.

Looking at classification practice, most AI diagnostic assistance software is placed in Class 3. Chest CT reading assistance, brain CT hemorrhage detection, mammogram reading, and retinal image diabetic retinopathy detection software fall into this category. Even if software assists rather than replaces a physician's judgment, it escalates to Class 3 if the reading result influences diagnosis. Class 4 applies to AI that makes independent diagnosis or treatment decisions without physician intervention, but there has not yet been a

case where a Class 4 AI medical device was approved in South Korea.

The approval review period was a problem. It took an average of 390 days. After the "Integrated Review and Assessment System for Innovative Medical Devices" was introduced in October 2022, it was reduced to 80 days. From the perspective of companies, market entry five times faster became possible.

Obtaining approval does not mean it can be used in hospitals immediately. Health insurance reimbursement rates must be set. As of July 2024, 8 AI products were designated as "Innovative Health Technology," and 7 received a "Deferral of New Health Technology Assessment." The reimbursement set under selective coverage ranges between 310 won and 2,920 won. With only a few thousand won added per CT reading, the revenue model for AI medical device companies remains unstable.

3. Digital Medical Products Act: The World's First Dedicated Law

In January 2024, the Digital Medical Products Act was promulgated. It created a separate law exclusively for digital medical products for the first time in the world.

This law distinguishes between three categories. Digital medical devices are medical devices to which AI and advanced digital technologies are applied. Digital combination drugs are medicines combined with digital medical devices or digital medical and health support devices. Digital medical and health support devices are digital technology products for healthcare support and health maintenance that do not constitute medical devices.

The enforcement dates of major provisions are divided into two stages. Most of the main provisions came into effect on January 24, 2025, and provisions related to digital medical and health support devices (Articles 33 through 35) came into effect on January 24, 2026.

The link to the AI Basic Act lies precisely here. Article 2, Item 4, Subitem (d) of the AI Basic Act lists medical devices pursuant to the Medical Devices Act side by side with "digital medical devices pursuant to the Digital Medical Products Act." The enforcement date for the digital medical devices portion of Subitem (d) is January 24, 2026. Since the main provisions of the Digital Medical Products Act defining digital medical devices had already been in effect since January 24, 2025, this date aligns with the day the Digital Medical Products Act completes full implementation through its final-stage

provisions.

Unpacking the relationship between the two laws yields the following: The Digital Medical Products Act falls under the jurisdiction of the MFDS, reviewing product safety and efficacy. The AI Basic Act falls under the jurisdiction of the Ministry of Science and ICT, managing the social impact and trustworthiness of AI systems. The same AI medical device is subject to simultaneous regulation by both ministries.

4. Healthcare AI as High-Impact Artificial Intelligence

When unpacking Subitems (c) and (d) of Article 2, Item 4 of the AI Basic Act, the scope under which healthcare AI is classified as high-impact is broader than expected.

Subitem (c)'s "establishment and operation of healthcare provision and usage systems" draws from Article 3, Item 1 of the Framework Act on Health and Medical Services. The Framework Act on Health and Medical Services defines healthcare as "all activities conducted by the State, local governments, healthcare institutions, or healthcare professionals, etc., to protect and promote the health of the public." It can encompass not only AI reading systems used in consultation rooms, but also hospital patient assignment algorithms, emergency

room triage AI, and infectious disease prediction systems.

Subitem (d) is more direct. It targets the "development and use" of medical devices pursuant to the Medical Devices Act and digital medical devices pursuant to the Digital Medical Products Act.

Business operators developing medical devices equipped with AI, as well as hospitals using those devices to provide services to patients, may fall within its coverage.

Professor Lee Won-bok (Ewha Womans University Law School) explicitly points out this aspect: "Not only medical devices and digital medical devices equipped with artificial intelligence, but also artificial intelligence systems used by healthcare institutions to facilitate medical care are automatically classified as high-impact artificial intelligence."

Once classified as high-impact AI, the obligations imposed on business operators are contained in five items under a single article, Article 34, Paragraph 1. Establishment and operation of risk management measures (Item 1). Establishment and implementation of explanation measures regarding final results, key criteria used to derive them, and an overview of training data (Item 2). Establishment and operation of user protection measures (Item 3). Human

management and supervision of high-impact artificial intelligence (Item 4). Preparation and retention of documents confirming the contents of measures taken (Item 5). Key details of the risk management measures, explanation measures, and user protection measures, along with the name and contact information of the supervising person, must be posted at the office or on the website, etc. (Enforcement Decree Article 27, Paragraph 1), and documents supporting the measures must be retained for 5 years (Enforcement Decree Article 27, Paragraph 2).

AI medical device development business operators approved by the MFDS have already submitted clinical data and technical documents during the approval process. Although there are overlaps with the documentation obligations under the AI Basic Act, the requirements of the laws do not completely align. MFDS technical documentation focuses on product safety and efficacy. Documentation under the AI Basic Act demands coverage of fundamental rights impact, bias management, and human intervention systems. The document names differ, the governing agencies differ, and the submission destinations differ.

5. Specific Cases: From Image Reading to Drug Discovery

The field with high maturity in the AI medical device market is image reading.

Lunit's "Lunit INSIGHT" is AI software that displays abnormal findings in chest X-rays and mammograms. It has received both MFDS approval and EU CE certification, and is also sold in South America and Southeast Asia. VUNO's "VUNOMed" is a product line analyzing chest X-ray, CT, and MRI images, with several products holding EU CE certification, and the company has also developed a cardiac arrest prediction algorithm. JLK supplies AI that performs early diagnosis and lesion detection from brain MRIs.

These companies are all development business operators under the AI Basic Act. They develop AI and provide it to hospitals. Operating in the high-impact AI domain (Subitem (d)), they bear the obligations of high-impact AI business operators.

Domains beyond image reading are also growing rapidly.

AI drug discovery uses machine learning for candidate compound screening and molecular structure optimization. It is the task of selecting molecules with a high likelihood of binding to specific protein targets from large compound libraries. Whether this AI system

constitutes a medical device under the Medical Devices Act is debated. If it is a research tool not directly used on patients, it may not be a medical device. However, if it is involved in healthcare provision and usage systems, the possibility remains that it could be classified as high-impact AI under Subitem (c).

AI genomic analysis applies AI to gene variant detection, expression pattern classification, and disease risk prediction. If genomic analysis software is approved as a medical device, Subitem (d) applies. Analysis tools used solely for research purposes have unclear applicability under the Medical Devices Act, but if their results influence clinical decision-making, they may be included in healthcare provision systems (Subitem (c)).

Sky Labs' ring-type blood pressure monitor "CART BP" received MFDS approval in 2023. It is the world's first cuffless blood pressure monitoring device to pass European Society of Hypertension (ESH) validation. Selvas AI's "Selvy MediVoice" automatically transcribes medical records based on speech recognition. Life Semantics' "CanopyMD SCAI" screens for skin cancer using smartphones. MediWhale evaluates cardiovascular risk from retinal images.

The question these products face in common is identical: To what extent do documents submitted to the MFDS overlap with, and diverge from, those required by the AI Basic Act?

6. Duplication of Obligations and Resolution: Limitations of Article 5

Article 5 (Relationship with Other Acts) of the AI Basic Act is a general provision intended to resolve this issue.

Paragraph 1 reads as follows: "Except as otherwise provided in special provisions of other Acts regarding artificial intelligence, etc., this Act shall govern." If there are "special provisions" related to AI in the Medical Devices Act or the Digital Medical Products Act, those provisions take precedence over the AI Basic Act.

The problem is that the scope of "special provisions" is ambiguous. The Medical Devices Act contains detailed provisions regarding the review of safety and efficacy for AI medical devices. It remains unclear whether these provisions constitute "special provisions" in relation to the obligation to establish risk management measures under Article 34, Paragraph 1, Item 1 of the AI Basic Act, or whether they are merely different regulations regarding different matters.

Professor Jeong Chan-mo (Inha University Law School) cites the precedent where a similar provision in Article 6, Paragraph 1 of the Personal Information Protection Act failed to resolve conflicts with other personal information-related laws in practice, pointing out that the general provision of Article 5 of the AI Basic Act alone can hardly resolve regulatory duplication with the Medical Devices Act or the Digital Medical Products Act.

Professor Lee Won-bok speaks more directly. He views that the general provision of Article 5 "can hardly be expected to substantially help resolve regulatory duplication between the AI Basic Act and the Medical Devices Act or the Digital Medical Products Act in the field of healthcare AI." Unless provisions in other laws qualify as "special provisions," business operators must comply with both laws simultaneously.

Compared with the European Union, the EU AI Act includes specific mechanisms to reduce duplication with the Medical Devices Regulation (MDR). Article 43, Paragraph 3 of the EU AI Act allows a single conformity assessment through an existing notified body. Article 11, Paragraph 2 allows a single technical document to meet the requirements of both laws. South Korea's AI Basic Act also has an

interlocking mechanism. Article 34, Paragraph 3 provides that where a business operator has implemented measures equivalent to each item of Article 34, Paragraph 1 pursuant to other statutes and regulations, such measures shall be deemed to have been implemented. Article 27, Paragraph 5 of the Enforcement Decree and Annex 1 specifically lay down recognition criteria. For digital medical devices, if a quality management system is established pursuant to the Digital Medical Products Act and a determination of conformity with quality management standards for digital medical device software is obtained, risk management measures (Item 1) and documentation (Item 5) are deemed implemented; if manufacturing/import approval, certification, or notification is completed and information is displayed on the software, explanation measures (Item 2) are deemed implemented; if all compliance matters under Article 13 of the same Act are fulfilled, user protection measures (Item 3) are deemed implemented; and if a quality manager is designated, human management and supervision (Item 4) are deemed implemented. However, these criteria are limited to digital medical devices under the Digital Medical Products Act. For traditional medical devices that received approval under the Medical Devices Act without going through the Digital Medical Products Act, there are no deeming criteria under Annex 1, leaving the limitations of

Article 5's general provision intact.

However, there is one realistic buffer. The maximum administrative fine under the AI Basic Act is 30 million won, and there are no criminal penalty provisions for violations of business operator obligations.

The government announced a policy to operate a guidance period of at least one year following the law's enforcement. Under this policy, no administrative fines will be imposed during that period, granting business operators time to reconcile obligations under both laws.

However, once the guidance period ends, the problem of duplicated obligations remains intact.

7. European Union: Dual Regulation under the MDR and the AI Act

In the European Union as well, healthcare AI stands at the intersection of two laws: the Medical Devices Regulation (MDR, Regulation (EU) 2017/745) and the AI Act (Regulation (EU) 2024/1689).

The EU AI Act has two pathways for classifying healthcare AI as high-risk. The Annex I pathway is where AI is included in a medical device requiring a third-party conformity assessment under the MDR or the In Vitro Diagnostic Medical Devices Regulation (IVDR). When AI is incorporated into a medical device of Class IIa or higher, it is automatically classified as high-risk AI. The Annex III pathway

designates designated uses as high-risk, such as emergency call patient classification and emergency triage, as well as eligibility evaluation for essential public services, capturing healthcare AI that falls under these categories.

Four points of friction arise from this dual structure.

First, conformity assessment procedures differ. The MDR certification procedure and the AI Act conformity assessment procedure (Article 43) differ in document formats, assessment criteria, and governing bodies. Although Article 43, Paragraph 3 of the EU AI Act permits a single assessment through an existing notified body, in practice the technical documentation format of the MDR (MDR Annex II) and that of the AI Act (AI Act Annex IV) differ in structure.

Second, a legal void exists regarding clinical investigations. The MDR and IVDR explicitly exclude investigational medical devices from the definitions of "placing on the market" or "putting into service." Pre-approval clinical investigations are permitted. The AI Act lacks a precisely corresponding exclusion provision, leading to interpretative debate over whether clinical investigations of AI medical devices constitute "putting into service" under the AI Act.

Third, risk classifications escalate. The MDR views Class IIb and above as true high-risk. However, the AI Act classifies devices from Class IIa onward as high-risk AI if third-party certification is required. An electronic thermometer equipped with AI may be low-risk under the MDR, but high-risk under the AI Act. According to a study published by Muehlematter et al. in *The Lancet Digital Health* in 2021, among 222 AI/ML medical devices approved by the US FDA, only 3 were in the highest risk tier (Class III PMA), and in the EU, only 3 out of 240 were Class III. The vast majority of the remainder were medium-to-low risk tiers, which could be reclassified as high-risk when the AI Act applies.

Fourth, deployer obligations are imposed on physicians.

8. Is a Physician a User or a Deployer?

The legal status of a physician differs dramatically between South Korea and the European Union.

In the European Union, a physician is a deployer. Article 3, Paragraph 4 of the EU AI Act defines a deployer as "any natural or legal person using an AI system under its authority." Only personal, non-professional activities are excluded. When a physician uses AI diagnostic assistance software in patient care, that physician (or

hospital) is a deployer. Deployers are subject to obligations under Article 26: using the AI in accordance with the manufacturer's instructions for use, monitoring risks, supervising operation, maintaining transparency, keeping logs, and complying with relevant laws and ethics. If the healthcare AI falls under Annex III, deployers that are public entities or private entities providing public services must even conduct a fundamental rights impact assessment pursuant to Article 27.

The situation is different in South Korea. Article 2, Item 8 of the AI Basic Act defines "user" as "a person who is provided with artificial intelligence products or artificial intelligence services." An "user business operator" under Article 2, Item 7, Subitem (b) is "a person who provides artificial intelligence products or artificial intelligence services by using artificial intelligence provided by a development business operator."

When a physician uses an AI medical device in patient care, is the physician a user business operator or a user? Interpretations diverge. Professor Lee Won-bok analyzes: "Depending on interpretation, medical professionals or healthcare institutions may or may not be viewed as falling under user business operators." If a hospital equips

itself with an AI reading system and provides image diagnostic services to patients, this constitutes an act of "providing services using artificial intelligence," which could classify it as a user business operator. On the other hand, if a physician uses AI merely as a tool in their primary duty of care, they could be viewed as "a person who is provided with" AI—namely, a user.

The practical consequences of this ambiguity are not minor.

If classified as a user, the business operator obligations under the AI Basic Act are exempted. Risk management measures, document retention, and website postings are not required. If classified as a user business operator, measures regarding risk management, explanation, and user protection that have already been implemented by the development business operator are deemed implemented unless there is a significant functional change (Enforcement Decree Article 27, Paragraph 3). However, if the original purpose or intended use has been significantly altered, independent obligations are borne.

In the European Union, this distinction is relatively clear. The physician is a deployer, and the patient is an affected person. All obligations under Article 26 of the EU AI Act apply to the physician as

deployer. Concerns are raised that this may act as a burden delaying AI adoption in healthcare settings. A paper published by van Kolfshoeten and van Oirschot in the journal Health Policy in 2024 analyzes the deterrent effect of deployer obligations under the EU AI Act on healthcare AI adoption.

The ambiguity in Korean law paradoxically provides flexibility. The Ministry of Science and ICT can clarify the status of healthcare institutions in either direction through guidelines or official interpretations. However, flexibility comes at the price of predictability. When a hospital considers adopting an AI medical device, if it cannot be certain of the obligations it will bear, postponing adoption itself becomes a rational choice.

9. What Business Operators Need to Verify Now

The checklist for AI medical device development business operators is as follows.

First, check whether MFDS approval has been obtained. If approved as an AI medical device, obligations under the Medical Devices Act are already being fulfilled. Differentiate between portions of high-impact AI obligations under the AI Basic Act that are satisfied by documents submitted to the MFDS and those that need to be prepared

additionally. If it is a digital medical device that has obtained approval, certification, or notification under the Digital Medical Products Act and possesses a quality management system and quality manager, first check whether a substantial portion of the obligations under each item of Article 34, Paragraph 1 is deemed implemented according to Annex 1 of the Enforcement Decree. The establishment of risk management measures (Article 34, Paragraph 1, Item 1) may substantially overlap with the MFDS risk management report, but human management and supervision systems (Article 34, Paragraph 1, Item 4) and impact assessment from the perspective of fundamental rights (Article 35) are highly unlikely to be included in MFDS documents.

Hospitals and healthcare institutions as user business operators must first determine their own legal status. The scope of obligations varies depending on whether they purchase AI medical devices as-is for use in care or provide new AI-based services to patients. They can consult the Ministry of Science and ICT's AI Basic Act Support Desk.

Business operators expanding into the EU market must prepare for dual regulation. Four regulatory layers pile up: South Korean MFDS approval, AI Basic Act obligations, EU MDR conformity assessment,

and EU AI Act high-risk AI obligations. Drafting technical documentation from the beginning in a format that satisfies the requirements of both laws can reduce redundant work.

The growth rate of South Korea's healthcare AI market is projected at 50.8% annually (from 2023 to 2030). This is faster than the global average of 41.8%. In a rapidly growing market, regulatory uncertainty becomes an obstacle to upfront investment. Business operators must keep monitoring whether amendments to the Enforcement Decree of the AI Basic Act or guidelines related to healthcare AI from the Ministry of Science and ICT can reduce this uncertainty.

Healthcare AI regulation is a structure that requires passing through two doors. If you assume you only need to pass through one door, you will be stopped in front of the remaining door. You must hold a map showing where the keys to the two doors overlap and where they diverge.

Chapter 29. Platforms and Generative AI Business Operators

March 2026, a startup office in Gangnam, Seoul. This company operates a service that automatically generates real estate listing descriptions by calling OpenAI's API. The CEO wrote three questions on the conference room whiteboard: Must we attach "Generated by AI" to the text we produce? Must that indication be visible text, or can it be hidden in metadata? If OpenAI has already fulfilled the indication obligation, do we not need to do so? The transparency obligation set forth in Article 31 of the AI Basic Act lands with different weight depending on each business model that uses generative AI.

1. The Weight of Obligations Varies Depending on the Business Model

Article 31 of the AI Basic Act stacks transparency obligations in three layers.

Paragraph 1 concerns prior notification. To provide a product or service using high-impact artificial intelligence or generative artificial intelligence, the fact that the product is operated based on the relevant artificial intelligence must be notified to users in advance. This applies commonly to all generative AI services.

Paragraph 2 concerns indication on generated output. An indication that "this was created by AI" must be attached to output produced by generative artificial intelligence. Whether text, image, video, or audio, if it is the output of generative AI, this obligation attaches to it.

Paragraph 3 is the special provision for deepfakes. When creating and providing outputs such as virtual audio, image, or video that are difficult to distinguish from reality, notice or indication must be provided in a manner that users can clearly perceive. The standard is higher than in Paragraph 2. The condition "in a manner that users can clearly perceive" has been added to "shall indicate."

Which of the three-layered obligations applies depends on what kind of output the business operator produces. A chatbot service generating text and a video generation service synthesizing human faces bear different obligations under the exact same article of the exact same act.

2. Chatbot Services: Prior Notification and Indication on Generated Output

For business operators running chatbot services such as ChatGPT, Claude, and Gemini, Paragraphs 1 and 2 of Article 31 apply.

Let us first look at the prior notification under Paragraph 1. Article 23(1) of the Enforcement Decree enumerates four methods of notification: directly stating it on the product or writing it in the terms of service, displaying it on the user's screen or terminal, posting it at the place where the service is provided, or other methods recognized by the Minister of Science and ICT. For a web-based chatbot, it suffices to include a notice string such as "This service is operated based on artificial intelligence" at the top of the conversation screen. Burying it solely within the terms of service is permitted under the statutory language, but transparency guidelines recommend a location that users can easily recognize.

The indication on generated output under Paragraph 2 has two methods. Article 23(2) of the Enforcement Decree places "methods recognizable by humans" and "machine-readable methods" side by side. Even if the machine-readable method is selected, one condition attaches: the fact that the output was generated by generative artificial intelligence must be provided at least once through notice text, voice, or other means. This means that hiding it in metadata alone is insufficient.

In practice, the way this obligation is fulfilled differs by service. ChatGPT constantly displays the text "ChatGPT can make mistakes" at the bottom of the conversation screen. Claude discloses that it is an AI assistant on its service screen and notice text. As long as they provide services to the Korean market, whether such indications satisfy the requirements of Article 31(2) of the AI Basic Act becomes a subject of practical legal judgment.

Most text outputs from chatbot services do not fall under the special provision for deepfakes in Article 31(3), because they are not "virtual audio, images, or videos that are difficult to distinguish from reality." Text is intrinsically a medium in which a formal distinction between real and virtual is impossible. Text written by a human and text written by AI cannot be distinguished by appearance alone. What Paragraph 3 targets are outputs that perceptually imitate reality—that is, photo-like images, video resembling real humans, and audio sounding like authentic voices.

There are exceptions. Article 23(4) of the Enforcement Decree exempts all or part of the obligation in three cases: when the fact that it is based on AI is obvious from the service name or screen, when it is used solely for the business operator's internal operations, or when

designated by the Minister of Science and ICT. Since the service name "ChatGPT" itself reveals the AI-based nature, there is room for it to qualify for the obviousness exception. However, for a business operator using the service who creates a chatbot with a separate name by calling its API, this exception is highly unlikely to apply.

3. Image Generation Services: Watermarking and Deepfakes

For image generation services such as Midjourney, DALL-E, and Stable Diffusion, Article 31(2) and (3) can apply simultaneously.

Paragraph 2 applies to all generative AI outputs. An image created by AI must bear an indication of that fact. The issue lies in the method. One can directly embed visible text such as "AI Generated" into the image output. This is a method recognizable by humans. Source information can also be embedded into the image metadata according to standards such as C2PA (Coalition for Content Provenance and Authenticity). This is a machine-readable method. Even if the machine-readable method is chosen, a notice must be provided at least once.

The point where Paragraph 3 additionally applies is images that are "difficult to distinguish from reality." Photo-grade images generated by models such as Midjourney v6 or above fall under this. Images of

individuals that are hard to distinguish from real photos, or landscape images that faithfully reproduce actual places, are examples. In such cases, going beyond the general indication obligation of Paragraph 2, notice or indication must be provided in "a manner that users can clearly perceive." Article 23(3) of the Enforcement Decree sets forth two specific criteria: users must be able to easily verify the content visually, auditorily, or through software; and the age and physical/social conditions of primary users must be taken into account.

In plain terms, it means that visible watermarks or labels must be attached to AI images that look like real photos. Relying solely on metadata makes it difficult to satisfy the requirement of "a manner that users can clearly perceive" under Paragraph 3.

DALL-E 3 automatically embeds C2PA metadata into generated images. Midjourney, as of 2026, is known not to have mandated metadata insertion. Since Stable Diffusion is an open-source model, distributors decide whether to include watermarking functionality. For business operators running these services in the Korean market or creating services based on them, combining visible indications alongside machine-readable methods is a choice that reduces legal

risk.

The latter part of Paragraph 3 includes consideration for artistic and creative expressions. If the relevant output constitutes or forms part of an artistic or creative expression, notice or indication may be provided "in a manner that does not impair display or enjoyment." This means that when exhibiting an AI-created artwork, rather than stamping a watermark directly over the artwork, it can be noted in exhibition placards or captions.

China's approach is more detailed. The Measures for Labeling AI-Generated and Synthesized Content (人工智能生成合成内容标识办法), which came into effect on September 1, 2025, imposes a dual labeling obligation: explicit labeling (i.e., labels visible to users) and implicit labeling (i.e., metadata or digital watermarks). Both must be performed. Mandatory national standard GB 45438-2025 regulates the technical specifications. Whereas Korean law allows choosing between the two, Chinese law mandates both.

4. Video Generation Services: The High Wall of Deepfake Labeling

Video generation services such as Sora, Runway, and Pika are core targets of Article 31(3).

Video carries a higher risk of deception than images. It moves, makes sound, and progresses through time. Videos that make real individuals appear to say things they never said, or videos showing events that never occurred—such outputs fit precisely into "virtual videos that are difficult to distinguish from reality."

What specifically does "a manner that users can clearly perceive" required by Paragraph 3 mean in video media? According to the criteria in Article 23(3) of the Enforcement Decree, users must be able to easily verify it visually or auditorily. One method is to continuously display an "AI Generated" watermark in a corner of the video. Another method is to insert subtitles or audio announcements such as "This video was generated by artificial intelligence" at the beginning and end of the video.

A practical difficulty arises here. As video is shared on social media, downloaded, edited, and reuploaded, visible watermarks may be cropped out or erased. Video metadata gets stripped when passing through platforms. There is no clear legal answer yet as to how the law addresses situations where a video whose indication vanished during distribution reaches users, even though the original business operator fulfilled the indication obligation.

The obligor under Article 31(3) of the Act is the "AI business operator." A development business operator operating a video generation service, or a user business operator operating a video production service based on that service, bears this obligation. General users who merely share generated videos are not subject to this obligation. Users under Article 2(8) are not the addressees of the law.

Article 50(4) of the EU AI Act takes this a step further. It imposes an obligation to disclose that fact even on deployers who use AI systems generating deepfakes in a professional capacity. While Korean law imposes obligations only on business operators, the EU expands it to users who deploy such systems for business or professional purposes.

5. Voice and Music Generation Services: Indication of Synthetic Voice

Synthetic voices and AI music created by services such as Suno AI, ElevenLabs, and AIVA correspond to the "audio" in Article 31(3).

ElevenLabs clones human voices. If a user inputs a few minutes of voice samples, it reads any text in that person's voice. If the output is difficult to distinguish from an actual human voice, it is subject to Paragraph 3. The fact that it is a synthetic voice must be conveyed in

a manner that users can clearly perceive.

Audio is a medium onto which visual watermarks cannot be attached. One method is inserting an announcement at the beginning or end of the audio file such as "This audio was synthesized by artificial intelligence." There is also technology that embeds watermarks in frequency bands inaudible to human ears within the acoustic signal. Since Article 23(3) of the Enforcement Decree recognizes verification using software, combining inaudible watermarks with audible announcements is a compliance method that reduces legal risk.

Music produced by Suno AI presents a slightly different issue. Whether a song composed by AI constitutes "virtual audio that is difficult to distinguish from reality" leaves room for interpretation. If it is a song imitating a real singer's voice, Paragraph 3 applies. If it is music composed with a fictional voice, the general indication obligation of Paragraph 2 applies, but it remains unclear whether it meets the "difficult to distinguish from reality" requirement of Paragraph 3. Judgments may vary depending on the degree to which it imitates the voice or performance style of a real person.

China's Provisions on the Administration of Deep Synthesis of Internet Information Services list six technical types: generation and editing of text; generation and editing of voice; generation and editing of non-speech audio such as music and ambient sound; generation and editing of biological features such as faces in images and videos; generation and editing of non-biological features in images and videos; and generation and editing of digital humans and virtual scenes. Voice is placed in a separate item, and non-speech audio such as music in another separate item. This contrasts with Article 31(3) of the Korean Act, which combines three media—"audio, image, or video"—into a single statutory provision. China split them by medium, whereas Korea embraced them in one.

6. API Providers and API-Utilizing Services: Who Bears the Labeling Responsibility?

Companies providing OpenAI API, Anthropic API, and Google's Gemini API are development business operators under Article 2, Subparagraph 7(a): "a person who develops and provides artificial intelligence." Companies that call these APIs to build their own services are user business operators under Article 2, Subparagraph 7(b): "a person who provides artificial intelligence products or artificial intelligence services using artificial intelligence provided by a business operator under Subparagraph (a)."

The subject of transparency obligations under Article 31 is the "AI business operator." Both development business operators and user business operators are AI business operators. Therefore, both bear transparency obligations.

API providers, who are development business operators, have obligations in two directions. When directly providing services to users (such as the ChatGPT web interface or Claude app), they directly fulfill prior notification and indication on generated output. When providing only APIs, providing technical means so that user business operators can fulfill indication obligations comes closer to the development business operator's duty. Examples include OpenAI embedding C2PA metadata into API responses by default or including provenance identification codes in generated outputs.

A user business operator's obligation arises within the context of its own service. For a fintech company operating a customer service chatbot by calling OpenAI's API, it must place prior notification text on the chatbot screen and indicate that answers produced by the chatbot were generated by AI. Even if the API provider has already performed metadata-level indication, the user business operator must separately perform notification and indication visible to users of

its own service.

Article 34(3) of the Act contains a provision to prevent double burden. Although limited to measures for securing safety and reliability by high-impact AI business operators, its intent is that if measures corresponding to each item of Article 34(1) have been implemented pursuant to other statutes, they shall be deemed implemented. However, the transparency obligation under Article 31 contains no such explicit provision preventing double burden. The specific criteria for allocating indication obligations between API providers and API-utilizing services remains a gap that must be filled by future guidelines or authoritative legal interpretations.

Article 50 of the EU AI Act addresses this issue more clearly. Providers of AI systems that generate synthetic audio, image, video, or text—including general-purpose AI systems—must mark the outputs of the system in a machine-readable format (Article 50(2)), and deployers of an AI system that generates deepfakes must disclose that fact (Article 50(4)). Tailored obligations are distributed to each stage of the value chain. Because Korean law imposes obligations on a single entity called the "AI business operator," the boundary of responsibility between API providers and API users is

relatively blurred.

7. The Dual Status of Platform Operators

Platform operators such as Naver, Kakao, and Google can be development business operators and user business operators at the same time.

When Naver provides a search AI service based on HyperCLOVA X developed in-house, it is a development business operator with respect to HyperCLOVA X. If the same Naver adopts OpenAI's model to operate a separate translation service, it is a user business operator with respect to that service. The two statuses are divided on a service-by-service basis. A single legal entity does not hold only one status.

The point where this dual status clashes in practice is the scope of transparency obligations.

Naver, as a development business operator, bears foundational duties for HyperCLOVA X: safety of the model itself, management of training data, and embedding of output indication technologies.

Naver, as a user business operator, fulfills transparency notifications tailored to its service context for external AI models it adopts.

There are more complex cases. This occurs when a platform develops its own AI model, provides that model to external developers via API, and simultaneously allows external developers to list services built with that API on the platform's marketplace. If Kakao develops KoGPT, opens its API, and an external developer provides a chatbot built with that API on KakaoTalk Channels, Kakao is the development business operator regarding KoGPT, and the external developer is the user business operator. Whether Kakao bears additional obligations as the operator of the distribution platform KakaoTalk Channels is an issue not explicitly specified by law. This contrasts with the EU AI Act, which designates distributors and product manufacturers as separate addressees of obligations.

Determining whether AI is high-impact is also divided by service. Naver's search AI service is highly unlikely to constitute high-impact artificial intelligence, because it does not directly fall under the eleven domains enumerated in Article 2, Subparagraph 4, such as energy supply, healthcare systems, or biometric identification for criminal investigation. However, if the same Naver develops a recruitment evaluation tool using AI, it falls under Article 2, Subparagraph 4(g): "decisions or evaluations that significantly affect individual rights and obligations, such as recruitment or loan

screening." As a high-impact AI business operator, obligations to secure safety and reliability under Article 34 are added: establishing risk management measures, implementing explanation measures, operating user protection measures, human management and supervision, and document creation and retention. Five obligations are loaded onto a single service.

There is a place where platform operators can seek consultation. The AI Basic Act Support Desk is jointly operated by the Ministry of Science and ICT and the Korea Software Industry Association (KOSA). As of the 10th week of implementation, it received 552 consultation requests. Inquiries regarding transparency indication obligations accounted for the largest portion at 51 percent. Trade secret protection and anonymous consultation are available.

8. China's Prior Registration Obligation: A Different Regulatory Philosophy

Under Korean law, there is no provision requiring registration with the government prior to launching a generative AI service. In China, there is.

China's Interim Measures for the Management of Generative Artificial Intelligence Services took effect on August 15, 2023. This regulation stipulates that to provide generative AI services with

"public opinion attributes or social mobilization capabilities" to the public within China, one must apply for a security assessment and complete algorithm filing (备案). Algorithm filing is a procedure reported within 10 business days after starting service provision under the Administrative Provisions on Algorithm Recommendation, whereas the launch filing for generative AI models is operated in practice to be completed prior to launching the service.

The registration status illustrates the scale. As of August 2024, more than 190 large language models were reported to have completed registration. Based on announcements by the Cyberspace Administration of China, 346 generative AI services were registered by late March 2025, and 611 by early November 2025.

Summarizing the difference between Korea and China in one sentence: Korea allows services to launch without prior registration or permit with the government; China requires fulfilling registration obligations before services can launch.

This difference is a difference in regulatory philosophy. Korea's AI Basic Act aims for ex-post self-regulation. Business operators examine on their own whether their AI constitutes high-impact artificial intelligence and, if necessary, request confirmation from the

Minister of Science and ICT (Article 33(1)). It is not prior permission. China combines prior registration with content control. The Cyberspace Administration of China (CAC) identifies the existence of services in advance and enforces ex-post.

The intensity of enforcement also differs. The upper limit of administrative fines under Korea's AI Basic Act is 30 million KRW. Article 43(1) imposes an administrative fine of up to 30 million KRW for violations of prior notification obligations, failure to designate a domestic agent, or non-compliance with suspension orders or corrective orders. China's Interim Measures for Generative AI invoke penalty provisions of the existing Cybersecurity Law, Data Security Law, and Personal Information Protection Law. For major violations related to personal information, fines of up to 50 million RMB (approx. 10 billion KRW), business suspension, and service termination orders are possible under the Personal Information Protection Law. In 2025, the Cyberspace Administration of China launched the 'Qinglang' special cleanup operation to crack down on unregistered generative AI apps and mini-programs.

There is one more thing Korean business operators must pay attention to in Chinese regulation. The Measures for Labeling

AI-Generated and Synthesized Content, which took effect on September 1, 2025, mandated both explicit labeling (labels visible to eyes) and implicit labeling (metadata or digital watermarks) simultaneously for all AI-generated content. Mandatory national standard GB 45438-2025 regulates technical specifications. This contrasts with Korean law offering options between "methods recognizable by humans" or "machine-readable methods." For Korean business operators providing or planning to provide generative AI services to the Chinese market, incorporating a dual labeling framework into design from the start is a prudent choice.

China is also clear regarding API provision. The Interim Measures for Generative AI state in the definition of provider (提供者) that it "includes cases of providing services through programmable interfaces (APIs), etc." API providers also bear registration obligations and content safety obligations as service providers. Korean law classifies API providers as development business operators, but does not impose a separate registration obligation on the provision of APIs itself.

The EU stands between Korea's ex-post self-regulation and China's prior registration control. While imposing ex-ante conformity

assessments on high-risk AI, the EU AI Act started with a soft law mechanism called Code of Practice for general-purpose AI models. The regulatory landscapes of these three jurisdictions move at different speeds.

9. What Business Operators Must Verify Right Now

Business operators using generative AI in their business must check four aspects regarding their service.

Whether they are a development business operator or a user business operator. If they directly develop or fine-tune an AI model to provide it, they are a development business operator. If they build their service by calling an external AI, they are a user business operator. They can also fall under both.

Whether the output of their service is text, image, video, or voice. For text generation services, Article 31(1) and (2) apply. For services generating images, videos, or voices that are difficult to distinguish from reality, the higher standard of Paragraph 3 is added.

Whether their service falls into the realm of high-impact artificial intelligence. If the service touches upon the eleven domains listed in Article 2, Subparagraph 4, obligations to secure safety and reliability

under Article 34 come into play.

If they are an overseas business operator, whether they meet the criteria for designating a domestic agent. The four criteria set by Article 29(1) of the Enforcement Decree: total revenue of 1 trillion KRW or more in the preceding year; revenue of 10 billion KRW or more in the AI service sector in the preceding year; an average of 1 million or more daily domestic users over the preceding 3 months; or a history of receiving an administrative fine for failing to comply with a suspension order or corrective order. If even one criterion is met, a domestic agent must be designated and reported to the Minister of Science and ICT.

Let us return to that real estate startup in the Gangnam office. This company is a user business operator. It generates text. Real estate listing description text is not "virtual audio, image, or video that is difficult to distinguish from reality." The prior notification under Article 31(1) and indication on generated output under Paragraph 2 are this company's obligations. It simply needs to notify users on the service screen that it operates on an AI basis, and attach an indication that it was created by AI to the generated text. Even if OpenAI fulfilled indication at the API level, this company must separately perform an

indication visible to users of its own service. What the law asks is not "Did someone do it?" but "Can the user know?"

Chapter 30. Public Sector Supply Business Operators

In the spring of 2026, the CEO of an AI startup in Guro Digital Complex, Seoul, opens KONEPS (Korea ON-line E-Procurement System) of the Public Procurement Service. Fifteen notices for AI-related service projects posted by government ministries and local governments are listed. The Ministry of Education has issued a call for bids to build an AI student evaluation assistance system, and the National Health Insurance Service has announced the upcoming introduction of an AI for reading health checkup images. This CEO's company produces a natural language processing-based automatic civil complaint classification system. It has a track record of a pilot implementation at one local government. However, an unfamiliar expression appears in the announcement: "In accordance with Article 16, Paragraph 3 of the AI Basic Act, priority consideration shall be given to artificial intelligence products or artificial intelligence services confirmed by the Minister of Science and ICT." Products that have not received this 'confirmation' may still participate in the bidding, but are excluded from priority consideration. If placed side by side with a competitor's product that has received confirmation, it falls behind. The CEO opens the Enforcement Decree.

1. Priority Consideration: A Public Procurement Pathway Created by Law

Article 16, Paragraph 3 of the AI Basic Act reads as follows: "When state agencies, etc., intend to purchase products or services necessary for performing their duties or issue service contracts, they shall give priority consideration to artificial intelligence products or artificial intelligence services prescribed by Presidential Decree." A proviso is attached: "Provided, That this shall not apply to cases where the utilization of artificial intelligence technology is inappropriate due to the nature of the duties."

This provision was added by an amendment act (Act No. 21311) promulgated on January 20, 2026, and entered into force on July 21, 2026. It was a provision not included in the original act. The National Assembly inserted it to promote the AI industry. During the legislative process, it was discussed under the title 'Promotion of AI Adoption in the Public Sector.'

"Shall give priority consideration" is a mandatory provision. It is not discretionary. When state agencies, local governments, and public institutions pursuant to Article 4 of the Act on the Management of Public Institutions purchase goods or contract out services, if an AI product or service is suitable for the duties, they must examine AI first. However, it is not a mandatory purchase obligation. After

examining, if they judge it unsuitable, they can choose another method.

Article 15, Paragraph 4 of the Enforcement Decree is the key. "Artificial intelligence products or artificial intelligence services prescribed by Presidential Decree" means "products or services confirmed by the Minister of Science and ICT." A product or service must receive confirmation from the Minister of Science and ICT to become subject to priority consideration. The criteria and procedures for confirmation shall be publicly notified by the Minister of Science and ICT in consultation with the Minister of the Interior and Safety (Enforcement Decree Article 15, Paragraph 5).

For a business operator seeking to supply AI to the public procurement market, this confirmation is a de facto qualification for entry. Even if it is not a mandatory purchase under the law, when a confirmed product competes with an unconfirmed product, the public institution official will prioritize the confirmed product. This is because the law specifies that they "shall give priority consideration." While it is possible to bid without confirmation, one can be at a disadvantage to the extent of being excluded from priority consideration.

2. Confirmation and Verification/Certification: Two Gateways

Before an AI supply business operator in the public sector, there are two gateways.

The first gateway is the confirmation by the Minister of Science and ICT. It is based on Article 15, Paragraph 4 of the Enforcement Decree. Once this confirmation is obtained, the product becomes subject to priority consideration in all public institution procurements. The criteria and procedures for confirmation are prescribed by public notice by the Minister of Science and ICT in consultation with the Minister of the Interior and Safety (Enforcement Decree Article 15, Paragraph 5). At the government announcement stage, a plan was proposed where relevant associations would handle application reception and confirmation issuance while Telecommunications Technology Association (TTA) would handle technical reviews; however, the finalized performing agency and detailed procedures must be verified through the public notice.

The Korea AI Promotion Association is established under Article 26 of the AI Basic Act. It is a legal entity established with authorization from the Minister of Science and ICT by persons engaged in AI-related research and operations. The Association's statutory duties include promoting the utilization of AI technologies and products, conducting

status surveys, nurturing specialized manpower, and supporting overseas expansion. If the Association takes on the role of determining technology application during the public procurement confirmation process, this Association becomes the gatekeeper of the procurement market for business operators.

The second gateway is verification and certification. Article 30 of the Act serves as its legal basis. To secure the safety and reliability of AI, the Minister of Science and ICT supports verification and certification activities autonomously pursued by the private sector. The support details include disseminating guidelines, supporting research, constructing equipment and systems, and nurturing specialized manpower. For small and medium-sized enterprises (SMEs), all or part of the verification and certification costs may be subsidized (Enforcement Decree Article 22, Paragraph 3).

The point where this verification and certification exerts power in the public sector is Article 30, Paragraph 4. "When state agencies, etc. intend to use high-impact artificial intelligence, they shall give priority consideration to products or services based on artificial intelligence that has received verification, certification, etc." To supply public institutions in the high-impact AI domain, products that

have received verification and certification take priority.

To summarize the relationship between the two gateways: To enter public procurement with a general AI product, confirmation from the Minister of Science and ICT is essential (Article 16, Paragraph 3). To supply a public institution with a high-impact AI product, verification and certification are added on top of that (Article 30, Paragraph 4). A business operator must first determine whether its product corresponds to high-impact AI, and if so, must also receive verification and certification to become subject to complete priority consideration.

Article 30, Paragraph 3 states that a business operator providing high-impact AI "shall endeavor to receive verification, certification, etc., in advance." It is a duty to endeavor. It is not mandatory. However, because Article 30, Paragraph 4 makes it mandatory for state agencies to give priority consideration to verified and certified products, the duty to endeavor becomes a de facto competitive requirement for business operators seeking to enter the public market. While the statutory text is gentle, the market logic is firm.

3. CAIOs of 48 Agencies: Where Purchasing Decisions Are Made

Following the implementation of the AI Basic Act, it is necessary to understand where public sector AI purchasing decisions are made. On September 18, 2025, Chief AI Officers (CAIOs) were designated across all 48 central administrative agencies (ministries, authorities, and agencies). The CAIO of a ministerial-level agency is assumed by a Vice Minister-level official, and the CAIO of a Vice Ministerial-level agency is assumed by a Deputy Minister/Director General-level official. On the same day, the CAIO Council was launched, which holds regular monthly meetings.

The CAIO oversees the AI adoption strategy of the relevant agency. They determine which duties will incorporate AI, what level of AI will be used, and how much budget will be allocated. For an AI business operator to enter the government procurement market, their product must ultimately align with the priorities of this CAIO.

The CAIO Council operates in connection with the National AI Strategy Committee. The Strategy Committee consists of up to 60 members in total, including 1 Chairperson (the President), 18 ex-officio members (heads of 16 central administrative agencies, the Deputy Director of the National Security Office handling AI duties, and the Senior Secretary to the President for AI duties), and up to 41

appointed civilian members, with appointed members required to constitute a majority of the total (Article 7, Paragraph 2 of the Act). The 3 Vice Chairpersons are not separate headcount allocations but are designated by the President from among ex-officio or appointed members, with 1 serving full-time. The Strategy Committee finalized the Republic of Korea Artificial Intelligence Action Plan (Basic Plan 2026–2028), a 199-page document containing 99 implementation tasks and 326 policy recommendations. The CAIO Council coordinates the agency-by-agency implementation of this Action Plan.

The structure that is practically important to AI business operators is as follows: The Strategy Committee sets the national-level AI policy direction, the CAIO Council conveys it to the 48 agencies, and the CAIO of each agency sets the AI adoption priorities for their own ministry. When the Public Procurement Service posts an announcement on KONEPS, business operators bid. At the top of this flow stands the Strategy Committee, while practical decisions are made by the CAIO.

There is a misstep made by companies attempting to win public sector AI contracts: focusing solely on the Public Procurement

Service. The Public Procurement Service is an entity that manages procurement procedures, not an entity that decides what to buy. What to buy is determined by the CAIO and operational departments of each ministry. Business operators must identify what AI projects the relevant ministry is planning and what direction is being discussed in the CAIO Council before a procurement announcement is issued.

On June 30, 2026, the AI Policy Center officially opened. The National Information Society Agency (NIA) was designated pursuant to Article 11 of the Act (NIA is the only entity designated to date), structured into 3 teams and 1 division, totaling 5 teams and 59 staff members. Its statutory projects include supporting the formulation of basic plans, supporting policy development, surveying and analyzing the impact of AI diffusion, and conducting legal and institutional research. This center serves as the practical hub for AI adoption in the public sector. Business operators can read the demand of the public sector through the adoption casebooks and technology status reports published by this center.

4. Immunity for Public Officials: The Key to Opening the Door

One of the reasons why selling AI to the public sector was difficult was the fear of public officials. If an AI system was introduced and

malfunctioned, triggering civil complaints, the official who approved the introduction was held accountable. They could be called out in an audit by the Board of Audit and Inspection and held liable for damages. The safe choice was not to introduce AI. If the existing method was maintained, nobody questioned it.

Article 16, Paragraph 4 of the AI Basic Act changes this structure: "Even if damage occurs to state agencies, etc., due to the purchase or use of artificial intelligence products or artificial intelligence services pursuant to Paragraph 3, the person responsible for such purchase or use duties shall not be liable to indemnify the agency for damages: Provided, That this shall not apply to cases where damage occurs due to intention or gross negligence of the person responsible for the relevant duties."

Let us unpack the structure of this provision. The subject of protection is "the person responsible for purchase or use duties." This refers to the public official who approves and operates the AI adoption. The scope of immunity is "liability for indemnification of damages to the relevant agency." This means that when monetary damage occurs to the agency due to AI, the official in charge does not have to pay that money out of pocket. The proviso clause is

"intention or gross negligence." Cases where a dangerous AI was introduced knowingly or where due care was extremely neglected are excluded from immunity.

There is a condition required for this immunity. It must be a product or service "pursuant to Paragraph 3." Immunity applies only when an AI product or service confirmed by the Minister of Science and ICT is purchased or used according to priority consideration procedures. If an unconfirmed AI product is arbitrarily purchased and a problem arises, it is difficult to be protected by this immunity.

Why is this immunity provision important from the perspective of business operators? When public officials feel less fear about AI adoption, purchasing decisions become faster. The law effectively provides an answer to the question, "What if a problem arises after incorporating AI?" If a product confirmed by the Minister of Science and ICT is properly adopted and there is no intention or gross negligence, the official in charge is not held responsible.

This structure also gives indirect protection to business operators. When public officials do not fear purchasing AI, the market opens up. When the market opens up, B2G revenue is generated. Although immunity is a provision for public officials, its effect connects to the

revenue of business operators.

However, this immunity waives the public official's liability for reimbursement to the agency, but does not waive the State's liability for state compensation to citizens. If damage occurs to a citizen due to a malfunction of an AI system, the State must compensate under the State Compensation Act. There is also room for exercising the right to indemnity against the business operator who developed the AI. Rather than being complacent because of the immunity clause, business operators must secure the safety of the product itself.

5. Item (i): When Public Institutions Make Decisions with AI

There is a provision that business operators supplying AI to the public sector must know. It is Item (i) of Subparagraph 4 of Article 2. Among the 11 domains of high-impact AI, this item directly applies to the public sector.

The original text of Item (i) reads as follows: "Decision-making by the State, local governments, public institutions pursuant to Article 4 of the Act on the Management of Public Institutions, etc., affecting citizens, such as verification and determination of qualifications necessary for providing public services or collection of fees."

Let us break down this sentence. The subjects are the State, local governments, and public institutions. The action is decision-making. The contents of decision-making are verification of qualifications, determination of qualifications, and collection of fees. Since "etc." is attached, these three are examples.

Applying concrete examples: Using AI for grading civil service examinations falls under Item (i), because it is qualification verification and determination. If AI determines eligibility for welfare benefits, it falls under Item (i). If AI is involved in tax assessment, it corresponds to fee collection. Cases where AI is used for calculating health insurance premiums, setting water utility rates, or making administrative decisions related to resident registration are also subsumed.

If an AI system corresponding to Item (i) is supplied to a public institution, that system becomes a high-impact AI. The responsibilities of a high-impact AI business operator follow. Pursuant to Article 34, establishing risk management measures, establishing explanation measures, user protection measures, securing human oversight and supervision systems, and document drafting and retention are required.

There is a point that business operators must be cautious of. The subject of Item (i) is public institutions, not business operators. Public institutions make decisions. However, the business operator who developed the AI system used for those decisions—that is, the developer business operator—bears the high-impact AI obligations. Depending on how it is used, a public institution can be a user or a business operator. Article 2, Subparagraph 7 includes state agencies, etc., in the definition of artificial intelligence business operators. When a public institution uses externally purchased AI for its administrative affairs, that institution becomes a user business operator.

As a result, both the developer business operator supplying AI to a public institution and the public institution using AI for administration (user business operator) become addressees of high-impact AI obligations. However, if the developer business operator has performed risk management measures, explanation measures, and user protection measures (Article 34, Paragraph 1, Subparagraphs 1 through 3), and if the public institution as the user business operator has not made significant functional changes, the obligation is deemed fulfilled to that extent (Enforcement Decree Article 27, Paragraph 3). Obligations for human oversight and

supervision, as well as document drafting and retention, must be separately performed by the public institution. This is a mechanism to reduce double burdens.

One more thing. If a public institution has already performed equivalent measures pursuant to other statutes as prescribed by Presidential Decree, it shall be deemed to have fulfilled the obligations under the Basic Act (Article 34, Paragraph 3; Enforcement Decree Article 27, Paragraph 5 and Attached Table 1). For example, if the National Tax Service uses AI in tax administration and complies with lawful procedures under the Framework Act on National Taxes or the Electronic Government Act, it may not need to separately fulfill the risk management obligations under the Basic Act.

6. Impact Assessment as a Weapon

In the public procurement market, impact assessment is both an obligation and a weapon.

Article 35 of the AI Basic Act allows artificial intelligence business operators to autonomously conduct high-impact AI impact assessments. It is not mandatory. However, business operators who perform this assessment secure an advantageous position in public procurement. This is the structure referred to as 'soft regulation'

during the legislative process. Although the law does not directly compel it, linking it with priority consideration in public procurement makes it operate as a de facto market entry condition.

The details to be included in the impact assessment are prescribed by Article 28 of the Enforcement Decree. There are seven items: identification of affected individuals and groups, identification of types of fundamental rights, social and economic impacts, usage patterns, evaluation metrics and calculation methods, prevention/mitigation of risks and loss recovery, and improvement implementation plans. Practical execution is guided by the Impact Assessment Guidelines across three stages: preliminary preparation, main assessment (writing risk scenarios and analyzing fundamental rights impacts), and follow-up measures. Business operators document and retain this assessment.

For a business operator seeking to supply AI to the public sector, completing an impact assessment is equivalent to a certificate stating, "Our product has examined impacts on fundamental rights and possesses a system to manage risks." A public institution official can present this document as evidence when undergoing an audit. In that it can document a rational reason for adopting AI, the impact

assessment also serves as a defense mechanism for public institutions.

The AI Impact Assessment Guidelines distributed by the Ministry of Science and ICT on January 22, 2026, set forth practical standards for this assessment. Business operators can conduct evaluations following these guidelines. Anonymous consultations are also available at the AI Basic Act Support Desk (operated by the Korea Software Industry Association). As of the 10th week of implementation, 552 consultation requests have been received.

7. National Defense and Security: The Market Outside the Law

Article 4, Paragraph 2 of the AI Basic Act is clear: "This Act shall not apply to artificial intelligence developed and used exclusively for national defense or national security purposes as prescribed by Presidential Decree."

The three words "exclusively for... purposes" are the key. Exemption applies only when it is dedicated to national defense or dedicated to national security. Dual-use technologies are not exempt. Even if an AI is supplied to the military, if it is a general-purpose technology that can also be used in the civilian sector, the Basic Act applies.

Article 2 of the Enforcement Decree narrows the scope of exemption. There are three designating entities.

Duties designated by the Minister of National Defense. These correspond to the construction and operation of defense information and communications networks and defense information systems, as well as the development and operation of weapon systems and military power support systems.

Duties designated by the Director of the National Intelligence Service. Enumerated tasks include operation of economic security early warning systems, security project research, construction of cybersecurity prevention and response systems, investigation of leaks of national advanced strategic technology, counter-terrorism activities, protection of defense industry technology, prevention of industrial technology leaks, electronic government security measures, integrated defense information tasks, counter-intelligence tasks, and duties under the National Intelligence Service Act.

Duties designated by the Commissioner General of the Korean National Police Agency. These correspond to investigation of advanced strategic technology leaks, counter-terrorism activities, investigation of defense industry technology leaks, investigation of

core technology leaks, counter-intelligence tasks, and investigation of information-related offenses.

This list is an exhaustive enumeration. Any national defense or security-related AI not listed here is subject to the Basic Act. Systems such as military meal management AI, military facility HVAC optimization AI, and soldier health management AI do not correspond to exemptions because they are not for combat or security purposes.

This distinction is vital for business operators seeking to enter the defense AI procurement market. If developing an AI subject to exemption, one is free from Basic Act obligations such as high-impact determination, impact assessment, and transparency duties. Instead, they become subject to defense-related laws such as the Defense Acquisition Program Act and the Military Secrets Protection Act. Exemption does not imply regulation immunity; it merely changes the governing law.

The EU AI Act has a similar structure. In Article 2, Paragraph 3, it exempts "AI systems placed on the market or put into service exclusively for military, defense or national security purposes." Both sides use the limiting qualifier "exclusively." The difference lies in the

fact that while the EU's statutory text itself is self-executing, South Korea delegates the specific scope to Presidential Decree. Article 2 of South Korea's Enforcement Decree exhaustively lists the duties. For duties not enumerated, it is difficult to claim an exemption.

8. Public AI Procurement in the European Union and China

Let us compare Korea's public sector AI priority consideration system with foreign countries.

The EU AI Act does not separately prescribe the promotion of AI adoption in the public sector. The EU's law is a regulatory law. Industrial promotion is pursued through separate policies. While the EU supports public sector AI adoption through the Digital Europe Programme, it did not include a provision in the AI Act itself stating that "public institutions shall give priority to purchasing AI." When public institutions in the EU adopt high-risk AI, they must mandatorily conduct a fundamental rights impact assessment. This is a regulatory provision, not a promotion provision.

South Korea blended promotion and regulation into a single law. Article 16, Paragraph 3 is promotion. Article 34 is regulation. From a business operator's perspective, entering the public procurement market is easier in Korea than in the EU. This is because the law

encourages public institutions to purchase AI. The EU's conformity assessment is a mandatory regulatory requirement that high-risk AI must pass regardless of whether it is procured or not. In contrast, Korea's confirmation and verification/certification are voluntary preferential mechanisms where receiving them brings the benefit of priority consideration. The nature of the threshold itself differs between the two systems.

China has no comprehensive provision for public sector AI adoption. Each ministry individually promotes AI pilot projects, while regulations are conducted by sector across nine or more ministries centered around the Cyberspace Administration of China (CAC). Article 17 of the Interim Measures for the Management of Generative Artificial Intelligence Services requires safety assessments and algorithm filings (备案) for generative AI services possessing public opinion attributes or social mobilization capabilities. Since most public-facing services fall under this category, it is difficult to provide such services without filing, whether in the public or private sector. As of November 2025, 611 generative AI services have been filed. To enter China's public sector AI market, this filing is a de facto prerequisite. Its nature differs from Korea's voluntary confirmation process.

Looking at the public sector AI procurement approaches of the three countries, clear differences emerge: South Korea says, "The law pushes the government to buy AI." The EU says, "If the government wants to buy AI, it must meet the safety requirements set by law." China says, "If an AI service can influence public opinion, whether government or private, it must be registered first."

9. Practicalities of Procurement Procedures

Let us follow the actual flow of public sector AI procurement.

KONEPS (G2B) of the Public Procurement Service is the central platform for public procurement. AI-related procurement is broadly divided into three pathways:

General bidding: The ordering agency sets the specifications, posts an announcement on KONEPS, and selects the successful bidder through qualification screening or technical evaluation. AI projects also follow this path. However, after the implementation of the AI Basic Act, bonus points or priority consideration may be awarded to products or services confirmed by the Minister of Science and ICT.

Multiple Award Schedule (MAS): A method where the Public Procurement Service enters into advance contracts with multiple

suppliers meeting certain quality standards, and demanding agencies select and purchase from among them. AI solutions can be listed under the software and information system categories.

Innovative Product Pilot Purchase Program: A system where the government pilot-purchases products that do not yet exist in the market or incorporate new technologies to test them in public institutions. This is a path that AI startups can use to take their first step into the public market.

The Ministry of Science and ICT annually operates the 'AI Voucher Program.' This program issues vouchers (coupons) so that SMEs can adopt AI solutions, through which AI supplier companies are selected. While public institutions do not directly adopt through this, it serves to verify the track record and capability of AI supplier companies. Track records accumulated in the voucher program serve as reference materials in public procurement bids.

10. Landscape of the B2G AI Market

The South Korean public sector AI market has acquired conditions for structural expansion following the implementation of the AI Basic Act.

There are legal demand-creation mechanisms: The priority consideration obligation under Article 16, Paragraph 3 encourages public institutions to purchase AI. The 48 CAIOs promote AI adoption in their respective agencies. Immunity for public officials lowers the psychological barrier to making adoption decisions.

Budget channels are open: Digital Platform Government implementation budgets are allocated to AI infrastructure. The Ministry of the Interior and Safety oversees the digital transformation of public institutions, while the Ministry of Science and ICT handles the diffusion of AI technology. Article 15, Paragraph 2 of the Enforcement Decree mandates that both ministries consult to establish concrete measures for supporting AI adoption.

There are cases of AI adoption already in operation: The National Tax Service's tax audit subject selection AI, the Korea Customs Service's imported goods search AI, the National Police Agency's CCTV analysis AI, Government24's civil complaint guidance AI chatbot, and the National Health Insurance Service's health checkup image reading assistance AI are already running in the field. These were introduced prior to the enforcement of the Act. Following the enforcement of the Act, re-examination is required to determine

whether these systems meet the obligations of the Basic Act, and the adoption of new systems will proceed toward equipping them with confirmation, verification, and certification.

However, just because the market opens up does not mean opportunities are equal for all AI business operators. Differences exist between large enterprises and startups in terms of Minister of Science and ICT confirmation procedures, verification and certification costs, and impact assessment execution capabilities. Article 30, Paragraph 2 of the Act provides for subsidizing verification and certification costs for SMEs, and Article 22, Paragraph 3 of the Enforcement Decree specifies that all or part of the costs may be subsidized within budget limits. The extent to which this support is actually provided will determine the entry of SME AI business operators into the public market.

To summarize public sector AI business in one sentence: Obtaining confirmation from the Minister of Science and ICT, determining whether it qualifies as high-impact, equipping it with verification, certification, and impact assessment if it does, and submitting a résumé to KONEPS. The person reading that résumé is the CAIO of 48 agencies.

There is one remaining question: The upper limit of the administrative fine under the AI Basic Act is 30 million KRW, and even that is imposed in only three cases: violation of prior notification, failure to designate a domestic agent, and failure to comply with a suspension or corrective order (Article 43, Paragraph 1). Violation of duties under Article 34 does not lead directly to an administrative fine, but goes through a fact-finding investigation and corrective order, leading to an administrative fine only when that order is not complied with. The cap on administrative fines (penalties) under the EU AI Act is 35 million EUR (approximately 50 billion KRW) or 7% of global annual turnover. However, the real cost in the public procurement market is not the administrative fine. If confirmation is revoked or trust is lost during the procurement process, a business operator can also be pushed out of procurements by other agencies. The public procurement market is a place where market exit is far more terrifying than financial sanctions.

Appendix

Appendix A. Business Operator Self-Diagnosis Flowchart

This flowchart has been prepared based on the AI Basic Act (enacted as Act No. 20676, amended by Act No. 21311, with amended provisions entering into force on July 21, 2026) and its Enforcement Decree (enacted as Presidential Decree No. 36053, reflecting amendments promulgated on July 20, 2026, and taking effect the following day). Business operators may follow it sequentially when determining the legal status and obligations of their services.

The determination consists of six steps.

Step 1. Does it qualify as an Artificial Intelligence System?

Basis Provision: Article 2, Subparagraph 2 of the Act

Question: Does your system satisfy all three of the following?

(a) Has autonomy and adaptability. Adjusts judgment criteria on its own based on input data.

(b) Derives outputs through inference. Performs one or more of prediction, recommendation, and decision-making.

(c) Influences physical or virtual environments.

Yes: Proceed to Step 2.

No: Not subject to the AI Basic Act. Ends here.

Note. Rule-based RPA, if-then macros, and software operating solely on fixed formulas lack autonomy and adaptability, and thus do not qualify as AI systems. Conversely, if a machine learning model serves as the core engine of the service, it is highly likely to qualify.

Step 2. Does it qualify as an AI Business Operator?

Basis Provision: Article 2, Subparagraph 7 of the Act

Question 2-1: Does your company directly develop AI and provide it to external parties?

Yes: You are a developer business operator. Proceed to Step 3.

Question 2-2: Do you provide products or services to external parties by utilizing AI developed by another business operator?

Yes: You are a deployer business operator. Proceed to Step 3.

Question 2-3: Do you only receive AI products or services without re-providing products or services to external parties based on them?

Yes: You are a user (Article 2, Subparagraph 8 of the Act). You do not bear business operator obligations under the Act. Ends here.

Note. An entity can hold both statuses simultaneously. If you develop AI in-house and also operate a loan screening service utilizing that AI, you are both a developer business operator and a deployer business operator. Corporations, organizations, individuals, and state agencies can all fall under this, regardless of for-profit status.

Note. If you fine-tune an external model and release it as a new service, this constitutes substantial development activity, and you may be classified as a developer business operator.

Step 3. Does it qualify as Generative AI?

Basis Provision: Article 2, Subparagraph 5 of the Act

Question: Does your AI system generate new outputs such as text, sound, images, or videos by imitating the structure and characteristics of input data?

Yes: It is generative AI. The following obligations are added:

Obligation (1) Output Labeling: You must label the outputs as being generated by AI in a manner that users can perceive or in a machine-readable format. (Article 31, Paragraph 2 of the Act; Article 23, Paragraph 2 of the Enforcement Decree)

Obligation (2) Prior Notification: You must notify users in advance that the product or service is operated based on artificial intelligence. (Article 31, Paragraph 1 of the Act; Article 23, Paragraph 1 of the Enforcement Decree)

Obligation (3) Deepfake Notification and Labeling: If you provide synthetic audio, images, videos, etc., that are difficult to distinguish from reality, you must notify or label them in a manner that users can clearly perceive. This is a separate obligation that applies regardless of generative AI status, and for artistic or creative expressions, it may be done in a way that does not impede exhibition or enjoyment. (Article 31, Paragraph 3 of the Act)

Proceed to Step 4.

No: It is not generative AI.

Labeling obligations related to generative AI do not apply, but separate obligations arise if it qualifies as high-impact AI or high-performance AI. Proceed to Step 4.

Note. Typical examples include systems that generate text, images, music, and videos, such as ChatGPT, Midjourney, and Suno AI. Systems that only perform classification or prediction and do not generate new content do not fall under this category.

Step 4. Does it qualify as High-Impact AI?

Basis Provisions: Article 2, Subparagraph 4 of the Act; Article 25 of the Enforcement Decree (Confirmation Procedures for High-Impact Artificial Intelligence, etc.)

The determination of high-impact AI consists of two sub-steps.

Step 4-1. Domain Verification

Is your AI system used in one or more of the following 11 domains?

- (a) Energy supply
- (b) Drinking water production process
- (c) Establishment and operation of healthcare provision and utilization systems
- (d) Development and use of medical devices and digital medical devices
- (e) Safety management and operation of nuclear materials and nuclear facilities
- (f) Analysis and utilization of biometric information (face, fingerprint, iris, palm vein, etc.) for criminal investigation or arrest
- (g) Decisions or evaluations that significantly affect individual rights and obligations, such as employment and loan screening
- (h) Major operation and management of transportation means, transportation facilities, and transportation systems
- (i) Decision-making by state agencies that affects citizens, such as qualification verification, determinations, and fee collection related

to public services

(j) Student evaluation in early childhood education, primary education, and secondary education

(k) Other domains prescribed by Presidential Decree

No: It is not high-impact AI. Proceed to Step 5.

Yes: Proceed to Step 4-2.

Step 4-2. Materiality Judgment

Does the AI system significantly affect or pose a potential risk to human life, physical safety, or fundamental rights?

No: It is not high-impact AI. Proceed to Step 5.

Yes: It is high-impact AI. The following obligations apply.

Major Obligations of High-Impact AI Business Operators:

Obligation (1) Establishment and Operation of Risk Management Measures: Must establish a system for risk identification, analysis, evaluation, treatment, and improvement. (Article 34, Paragraph 1,

Subparagraph 1 of the Act)

Obligation (2) Establishment and Implementation of Explanation

Measures: Must prepare explanation measures regarding final results, major criteria, and an outline of training data, and publish key contents at business offices, websites, etc. (Article 34, Paragraph 1, Subparagraph 2 of the Act; Article 27, Paragraph 1 of the Enforcement Decree)

Obligation (3) Establishment and Operation of User Protection

Measures (Article 34, Paragraph 1, Subparagraph 3 of the Act)

Obligation (4) Human Management and Supervision: Must provide

mechanisms allowing human intervention in AI decisions. (Article 34, Paragraph 1, Subparagraph 4 of the Act)

Obligation (5) Documentation: Must retain compliance

documentation for 5 years. (Article 34, Paragraph 1, Subparagraph 5 of the Act; Article 27, Paragraph 2 of the Enforcement Decree)

Obligation (6) Impact Assessment: Must endeavor to autonomously

assess impacts on fundamental rights. (Article 35 of the Act)

Proceed to Step 5.

Note. Confirmation on whether a system qualifies as high-impact AI may be requested from the Minister of Science and ICT. The confirmation period is 30 days, which may be extended once by 30 days. (Article 33 of the Act; Article 25, Paragraph 3 of the Enforcement Decree)

Step 5. Does it qualify as High-Performance AI?

Basis Provisions: Article 32 of the Act; Article 24 of the Enforcement Decree

Question: Does your AI system meet all three of the following?

(a) Cumulative computation used for training is 10^{26} floating-point operations (FLOPs) or greater.

(b) Configured and operated by applying AI technology corresponding to the current state of the art.

(c) Has the potential to exert widespread and significant impacts on human life, physical safety, and fundamental rights.

Yes: It is high-performance AI. The following obligations apply:

Obligation (1) Identification, evaluation, and mitigation of risks throughout the entire lifecycle (Article 32, Paragraph 1, Subparagraph 1 of the Act)

Obligation (2) Establishment of a risk management system for safety incident monitoring and response (Article 32, Paragraph 1, Subparagraph 2 of the Act)

Obligation (3) Submission of implementation results to the Minister of Science and ICT (Article 32, Paragraph 2 of the Act)

Proceed to Step 6.

No: It is not high-performance AI. Proceed to Step 6.

Note. All three requirements must be met. If only the amount of computation is large but the remaining requirements are not met, it does not qualify as high-performance AI.

Comparison: The systemic risk threshold for General-Purpose AI (GPAI) under the EU AI Act is 10^{25} FLOPs. South Korea's threshold is 10 times higher, so most domestic business operators will not fall into this category.

Step 6. Does it fall under the Obligation to Designate a Domestic Agent?

Basis Provisions: Article 36 of the Act; Article 29 of the Enforcement Decree

This step applies only to overseas business operators that do not have an address or place of business in South Korea.

If you have an address or place of business in South Korea, skip this step.

Question: Do any of the following four apply?

- (a) Total revenue for the preceding year is 1 trillion KRW or more.
- (b) Revenue in the AI service sector for the preceding year is 10 billion KRW or more.
- (c) Average daily number of domestic users during the preceding 3 months is 1 million or more.
- (d) Has been imposed an administrative fine under Article 43, Paragraph 1, Subparagraph 3 of the Act.

Yes: You must designate a domestic agent in writing and report to the Minister of Science and ICT.

The domestic agent must be a person who has an address or place of business in South Korea.

Matters represented by the domestic agent:

Submission of implementation results for high-performance AI (Article 32, Paragraph 2 of the Act)

Request for confirmation on whether an AI system qualifies as high-impact AI (Article 33, Paragraph 1 of the Act)

Support for the implementation of measures to ensure safety and reliability (Article 34, Paragraph 1 of the Act)

In case of failure to designate: Administrative fine not exceeding 30 million KRW (Article 43 of the Act)

No: No obligation to designate a domestic agent.

Note. Pursuant to Article 29, Paragraph 2 of the Enforcement Decree, an overseas business operator's revenue is based on the

amount converted into KRW using the average exchange rate of the preceding year.

Summary Table of Self-Diagnosis Results

The table below summarizes the sets of obligations based on the determination results. Please check the applicable box.

General AI Business Operators (Confirmed in Step 2)

If it does not fall under high-impact, generative, or high-performance AI, there are no individual obligations imposed by this Act. The prior notification obligation (Article 31, Paragraph 1 of the Act) arises when providing products or services utilizing high-impact or generative artificial intelligence.

Generative AI Business Operators (Confirmed in Step 3)

Prior notification obligation (Article 31, Paragraph 1 of the Act)

Output labeling obligation (Article 31, Paragraph 2 of the Act)

Deepfake notification and labeling obligation: When providing corresponding outputs (Article 31, Paragraph 3 of the Act)

High-Impact AI Business Operators (Confirmed in Step 4)

Prior notification obligation (Article 31, Paragraph 1 of the Act)

Establishment and operation of risk management measures (Article 34, Paragraph 1, Subparagraph 1 of the Act)

Establishment and implementation of explanation measures (Article 34, Paragraph 1, Subparagraph 2 of the Act)

Establishment and operation of user protection measures (Article 34, Paragraph 1, Subparagraph 3 of the Act)

Human management and supervision (Article 34, Paragraph 1, Subparagraph 4 of the Act)

5-year retention of compliance documentation (Article 34, Paragraph 1, Subparagraph 5 of the Act; Article 27, Paragraph 2 of the Enforcement Decree)

Efforts toward fundamental rights impact assessment (Article 35 of the Act)

High-Performance AI Business Operators (Confirmed in Step 5)

Risk identification, evaluation, and mitigation (Article 32, Paragraph 1, Subparagraph 1 of the Act)

Establishment of risk management system (Article 32, Paragraph 1, Subparagraph 2 of the Act)

Submission of implementation results to MSIT (Article 32, Paragraph 2 of the Act)

Overseas Business Operators' Domestic Agent (Confirmed in Step 6)

Written designation and reporting of domestic agent (Article 36 of the Act)

This flowchart reflects the current law as of July 21, 2026 (reflecting amended Act No. 21311) and the amended Enforcement Decree that took effect on the same date. Anonymous consultations are available at the Ministry of Science and ICT AI Basic Act Support Desk (www.sw.or.kr/AI_act_helpdesk).

Appendix B. Mapping Table of Statutory Provisions, Enforcement Decree, and Guidelines

Places core provisions of the AI Basic Act (enacted as Act No. 20676 and taking effect on January 22, 2026; amended by Act No. 21311,

with amended provisions taking effect on July 21, 2026) in Column 1, provisions of the Enforcement Decree detailing those articles (enacted as Presidential Decree No. 36053, reflecting amendments promulgated on July 20, 2026, and taking effect the following day) in Column 2, and relevant MSIT guidelines (5 major guidelines and 2 public notices distributed on January 22, 2026) in Column 3.

Abbreviations for the 5 Major Guidelines:

[Transparency] Guidelines for Ensuring AI Transparency

[Safety] Guidelines for Ensuring AI Safety

[High-Impact Judgment] Guidelines for Determining High-Impact AI

[Operator Duties] Guidelines for Duties of High-Impact AI Business Operators

[Impact Assessment] Guidelines for AI Impact Assessment

Abbreviations for the 2 Public Notices:

[Frontier Notice] Public Notice on Ensuring Safety of Frontier Artificial Intelligence

[Duties Notice] Public Notice on Duties of High-Impact AI Business Operators

Note: The article numbers in this table follow the current legal numbers verified on the National Law Information Center as of July 21, 2026. Draft bills from the National Assembly review stage or some explanatory materials retain different numbers (Article 15 Prior Notification, Article 16 Output Labeling, etc.), so the correspondence between old and current numbers is summarized in Remark 1.

Article (Basic Act) | Enforcement Decree (Current, Effective 2026.7.21) | Relevant Guidelines/Notices

=====

=====

Article 2, Subparagraph 1 (Artificial Intelligence)

"The realization of human intellectual abilities, such as learning, reasoning, perception, judgment, and language comprehension, through electronic methods" | No separate detail provision in the Enforcement Decree. | Cited as standard for determining AI qualification in [High-Impact Judgment].

The Act itself provides a self-contained definition. |

Article 2, Subparagraph 2 (Artificial Intelligence System)

A system that infers outputs such as predictions, recommendations, decisions, etc., with varying levels of autonomy and adaptability | No separate detail provision in the Enforcement Decree. | Explains autonomy and adaptability requirements with examples when judging system qualification in [High-Impact Judgment].

Rule-based automation (RPA) is not subsumed under this definition.

|

Article 2, Subparagraph 4 (High-Impact Artificial Intelligence)

An AI system that significantly affects or poses a potential risk to human life, physical safety, and fundamental rights. Enumerates 11 domains (items a through k). | Article 25 of Enforcement Decree: Confirmation procedure for high-impact status. Submission of confirmation request, reply within 30 days (+30 days extension), re-confirmation request submitted within 10 days. | Entirety of [High-Impact Judgment]. Elaborates 2-step determination (domain applicability + materiality) with detailed criteria and case examples.

Article 26 of Enforcement Decree: Establishment of a Special Committee. |

Article 2, Subparagraph 5 (Generative Artificial Intelligence)

A system that generates outputs such as text, sound, images, video, etc., by imitating the structure and characteristics of input data |

Article 23, Paragraph 2 of Enforcement Decree: Output labeling methods (human-perceivable method or machine-readable method).

| Guidance on output labeling methods for generative AI in [Transparency].

Article 2, Subparagraph 7 (AI Business Operator)

a. Developer Business Operator: A person who develops and provides artificial intelligence | Article 27, Paragraph 3 of Enforcement Decree: If developer business operator implements Article 34, Paragraph 1, Subparagraphs 1–3 measures, deployer business operator is deemed to have implemented them to that extent unless making significant functional changes. | Explains boundary of obligations between developer business operators and deployer business operators in [Operator Duties]. Includes cases of holding both statuses simultaneously.

b. Deployer Business Operator: A person who provides products or services using AI provided by developer business operator | Article 27, Paragraph 4 of Enforcement Decree: Deployer business operator may request materials from developer business operator. |

High-Performance Artificial Intelligence

(No definition in Article 2 of Basic Act. Prescribed in Enforcement Decree under mandate of Article 32 of Act) | Article 24 of Enforcement Decree: Qualifies if all three criteria are met. | Detailed guidance in [Safety] and [Frontier Notice] on compute calculation methods, risk identification/evaluation/mitigation procedures, and result submission methods.

- 1) Cumulative compute of 10^{26} FLOPs or more |
- 2) Application of state-of-the-art technology |
- 3) Concern of widespread and significant impacts on life/body/fundamental rights |

Appendix C. Checklists by Business Operator Type

Based on the AI Basic Act (enacted as Act No. 20676, amended by Act No. 21311) and its Enforcement Decree (enacted as Presidential

Decree No. 36053, reflecting amendments taking effect on 2026.7.21). As of July 21, 2026. Statutory basis provisions are indicated in parentheses for each item. The requirements for designating a domestic agent, which include four criteria, are in Article 29, Paragraph 1 of the Enforcement Decree. The 5 major guidelines are the guidelines on ensuring transparency, ensuring safety, determining high-impact AI, business operator duties, and impact assessment distributed by the Ministry of Science and ICT on January 22, 2026.

1. Checklist for Developer Business Operators

[] Verify whether company products or services meet the definition of AI system (Article 2, Subparagraph 2 of the Act)

[] Check compliance with autonomy, adaptability, and output derivation through inference requirements (Article 2, Subparagraph 2 of the Act)

[] Confirm that rule-based automation (RPA, macros) does not qualify as an AI system (Article 2, Subparagraph 2 of the Act)

[] Confirm whether company qualifies as a developer business operator: A person who develops and provides artificial intelligence

(Article 2, Subparagraph 7, Item a of the Act)

[] Check possibility of classification as developer business operator if providing a new service by fine-tuning external models (Article 2, Subparagraph 7, Item a of the Act)

[] Confirm qualification as generative AI: Whether the system generates outputs such as text, sound, images, video, etc. (Article 2, Subparagraph 5 of the Act)

[] Confirm qualification as high-impact AI, cross-referencing the 11 major domains (energy, drinking water, healthcare, medical devices, nuclear energy, biometric criminal investigation, rights/duties judgments such as recruitment and loan screening, transportation, public services, student evaluation, additional domains under Presidential Decree) (Article 2, Subparagraph 4, Items a through k of the Act)

[] Conduct Step 2 materiality judgment for high-impact AI: Assess whether it significantly affects life, physical safety, or fundamental rights (Article 25, Paragraph 2 of the Enforcement Decree; High-Impact AI Determination Guidelines)

[] Record grounds for high-impact determination in writing (Article 25 of the Enforcement Decree)

[] Confirm qualification as high-performance AI: Cumulative training compute of 10^{26} FLOPs or more (Article 32 of the Act; Article 24 of the Enforcement Decree)

[] Establish a system to notify users in advance that products or services are operated by AI when providing high-impact or generative AI products or services: Timing, means, and content of notification (Article 31, Paragraph 1 of the Act; Article 23, Paragraph 1 of the Enforcement Decree)

[] Attach visible labels to generative AI outputs (Article 31, Paragraph 2 of the Act; Article 23, Paragraph 2 of the Enforcement Decree)

[] Embed machine-readable labels (watermarks, etc.) in generative AI outputs (Article 31, Paragraph 2 of the Act; Article 23, Paragraph 2 of the Enforcement Decree)

[] Separately attach visible labels to outputs qualifying as deepfakes; invisible watermarking alone is insufficient (Article 31,

Paragraph 3 of the Act; Article 23, Paragraph 3 of the Enforcement Decree)

[] Review applicability of labeling exceptions if obviously AI or for internal business use only (Article 23, Paragraph 4 of the Enforcement Decree)

[] For high-impact AI, establish and operate risk management measures: Risk identification, analysis, evaluation, treatment, and improvement procedures (Article 34 of the Act; Public Notice on Duties of High-Impact AI Business Operators)

[] For high-impact AI, publish key criteria, training data outline, user protection measures, and management contact info at business offices, workplaces, or website (Article 34 of the Act; Article 27, Paragraph 1 of the Enforcement Decree; Public Notice on Duties of High-Impact AI Business Operators)

[] For high-impact AI, establish procedures to respond to explanation requests from affected persons: Explain major criteria and principles used to derive final results (Article 3, Paragraph 2 and Article 34 of the Act)

[] For high-impact AI, establish a human intervention system:
Structure where humans can intervene in or override AI decisions
(Article 34, Paragraph 1, Subparagraph 4 of the Act)

[] Establish a system to retain compliance documentation for 5
years (Article 34 of the Act; Article 27, Paragraph 2 of the Enforcement
Decree)

[] For high-performance AI, perform risk identification, evaluation,
and mitigation measures, and submit implementation results to MSIT
(Article 32, Paragraph 2 of the Act; Article 24 of the Enforcement
Decree)

[] Autonomously conduct fundamental rights impact assessment:
Prior preparation, main assessment (drafting risk scenarios and
analyzing fundamental rights impacts), post-treatment (Article 35 of
the Act; Article 28 of the Enforcement Decree; Impact Assessment
Guidelines)

[] Check training data for copyright infringement and log sources
(Copyright Act)

☐ Verify legality of processing personal information collected during training (Personal Information Protection Act)

☐ Organize information to be provided to deployer business operators: Characteristics, limitations, conditions of use, and risk factors of the AI system (Article 34 of the Act; Business Operator Duties Guidelines)

☐ Check eligibility for deemed compliance if equivalent measures have been implemented under existing statutes (Medical Devices Act, Use and Protection of Credit Information Act, Telecommunications Business Act, etc.) (Article 34, Paragraph 3 of the Act; Article 27, Paragraph 5 and Attached Table 1 of the Enforcement Decree)

☐ Check 5 major guidelines and verify reflection in internal procedures (MSIT Guidelines)

2. Checklist for Deployer Business Operators

☐ Confirm whether company qualifies as a deployer business operator: A person who provides products or services using AI provided by developer business operator (Article 2, Subparagraph 7, Item b of the Act)

[] Check distinction from user (Article 2, Subparagraph 8 of the Act):
If receiving AI merely to use in original business, entity is a user and not subject to obligations (Article 2, Subparagraph 8 of the Act)

[] Verify whether information regarding AI system characteristics, limitations, and conditions of use was received in writing from developer business operator (Article 34 of the Act; Business Operator Duties Guidelines)

[] Review received information in context of company's service (Business Operator Duties Guidelines)

[] For services using high-impact or generative AI, notify users in advance that service is operated based on AI (Article 31, Paragraph 1 of the Act; Article 23, Paragraph 1 of the Enforcement Decree)

[] Fulfill labeling obligations when providing generative AI outputs through company services (Article 31, Paragraph 2 of the Act; Article 23, Paragraph 2 of the Enforcement Decree)

[] Check whether company service falls under 11 major high-impact AI domains (Article 2, Subparagraph 4, Items a through k of the Act)

[] If falling under high-impact domains, separately establish risk management measures (Article 34 of the Act)

[] If developer business operator has already implemented risk management, explanation, and user protection measures (Article 34, Paragraph 1, Subparagraphs 1–3 of the Act) and company has not made significant functional changes, check whether conditions for deemed compliance are met within that scope; human management/supervision and documentation must be performed separately (Article 34 of the Act; Article 27, Paragraph 3 of the Enforcement Decree)

[] Recognize that if functionality is substantially modified, company assumes developer business operator obligations (Article 2, Subparagraph 7 of the Act; Business Operator Duties Guidelines)

[] Establish procedures to respond to explanation requests from affected persons (Article 3, Paragraph 2 of the Act)

[] Establish human intervention system when using high-impact AI (Article 34, Paragraph 1, Subparagraph 4 of the Act)

☐ Check personal information processing consignment relationship: Distinguish whether entity is consignor, consignee, or joint controller in relation to developer business operator (Article 26 of the Personal Information Protection Act)

☐ Establish user complaint reception and damage relief procedures (Article 3 of the Act)

☐ Reflect AI usage facts and limitations in Terms of Service (Transparency Guidelines)

☐ Establish a system to retain compliance documentation for 5 years (Article 34 of the Act; Article 27, Paragraph 2 of the Enforcement Decree)

☐ Check eligibility for deemed compliance if equivalent measures have been implemented under existing statutes (Article 34, Paragraph 3 of the Act; Article 27, Paragraph 5 and Attached Table 1 of the Enforcement Decree)

☐ Check 5 major guidelines and verify reflection in internal procedures (MSIT Guidelines)

3. Checklist for Overseas Business Operators

[] Check applicability of extraterritorial application: Whether company's AI service impacts the Korean market or Korean users (Article 4, Paragraph 1 of the Act)

[] Check extraterritorial nexus factors such as providing Korean language services, KRW payments, marketing targeting Korean consumers (Article 4, Paragraph 1 of the Act)

[] Check whether company has an address or place of business in Korea: If so, exempt from obligation to designate domestic agent (Article 36, Paragraph 1 of the Act)

[] If company has a Korean subsidiary, check whether contracting party for AI service is Korean entity or overseas headquarters (Article 36 of the Act)

[] Domestic agent designation criterion 1: Preceding year total revenue of 1 trillion KRW or more (Article 29, Paragraph 1, Subparagraph 1 of the Enforcement Decree)

[] Domestic agent designation criterion 2: Preceding year revenue in AI service sector of 10 billion KRW or more (Article 29, Paragraph 1, Subparagraph 2 of the Enforcement Decree)

[] Domestic agent designation criterion 3: Average daily domestic users during preceding 3 months as of end of preceding year of 1 million or more (Article 29, Paragraph 1, Subparagraph 3 of the Enforcement Decree)

[] Domestic agent designation criterion 4: If entity has history of administrative fines under Article 43, Paragraph 1, Subparagraph 3 of the Act, obligation arises regardless of size (Article 29, Paragraph 1, Subparagraph 4 of the Enforcement Decree)

[] If any one of four criteria applies, obligation to designate domestic agent arises (Article 29, Paragraph 1 of the Enforcement Decree)

[] Was domestic agent designated in writing and reported to Minister of Science and ICT (Article 36, Paragraph 1 of the Act)

[] Verify whether domestic agent is a person with an address or place of business in Korea (Article 36, Paragraph 2 of the Act)

[] Confirm delegation of three statutory roles of domestic agent: Submission of high-performance AI implementation results (Article 32, Paragraph 2 of the Act), responding to high-impact AI

confirmation requests (Article 33, Paragraph 1 of the Act), supporting implementation of safety and reliability measures for high-impact AI (Article 34, Paragraph 1 of the Act) (Article 36, Paragraph 1 of the Act)

[] Recognize that violations by domestic agent are attributed to overseas headquarters (Article 36, Paragraph 3 of the Act)

[] Provide prior notification to Korean users in Korean: AI operation facts, features, and limitations (Article 31, Paragraph 1 of the Act; Article 23, Paragraph 1 of the Enforcement Decree)

[] Apply output labeling obligations to services targeting Korean users (Article 31, Paragraph 2 of the Act; Article 23, Paragraph 2 of the Enforcement Decree)

[] Establish system to respond to high-impact AI confirmation requests (Article 33 of the Act)

[] For high-performance AI, establish system to submit safety assurance implementation results (Article 32, Paragraph 2 of the Act)

[] Verify applicability of Korean Personal Information Protection Act and establish compliance system (Personal Information Protection Act)

[] Reflect AI usage facts, limitations, and damage relief procedures in Terms of Service targeting Korean users (Article 31 of the Act; Transparency Guidelines)

[] Prepare response system in case of administrative fine:
Administrative fine up to 30 million KRW for failure to designate domestic agent (Article 43 of the Act)

[] Check eligibility for defense or national security exclusive exception: Mixed uses that are not exclusive cannot be exempted (Article 4, Paragraph 2 of the Act; Article 2 of the Enforcement Decree)

Appendix D. Implementation Schedule Timeline

Summarizes the implementation schedule of the AI Basic Act and related laws and regulations in chronological order.

2022.3.1 China Provisions on the Administration of Algorithmic Recommendations in Internet Information Services Take Effect

Business operators operating algorithm recommendation services in China must register (file/备案) algorithm types and application domains with the Cyberspace Administration of China. An obligation was created to provide users with the option to turn off algorithmic

recommendations.

2023.8.15 China Interim Measures for the Management of Generative Artificial Intelligence Services Take Effect

Business operators providing generative AI services within the territory of China bear obligations for prior algorithm registration (filing/备案) and content safety. This includes cases where services are provided via API.

2024.12.26 AI Basic Act Passes National Assembly Plenary Session (Act No. 20676)

Approved with 260 votes out of 264 present. Following the EU, the world's second comprehensive horizontal AI law was established. Immediate legal obligations do not yet arise for business operators, but this marks the point to begin preparations during the one-year grace period after promulgation.

2025.1.21 AI Basic Act Promulgated

The official title is the Basic Act on the Development of Artificial Intelligence and Establishment of Trust. Enforcement was confirmed for January 22, 2026, one year from promulgation. Business operators

should begin inspecting whether their services fall under the scope of this Act.

2025.2.2 EU AI Act Prohibited Practices Provisions Become Applicable

Prohibited practice provisions (Article 5), such as social scoring and real-time remote biometric identification, as well as AI literacy obligations, became applicable from this date. Korean business operators providing services to the EU market are also subject to these provisions.

2025.8.2 EU AI Act General-Purpose AI (GPAI) Obligations Become Applicable

Obligations of general-purpose AI model providers, such as drawing up technical documentation, establishing copyright policies, and publishing summaries of training data, became applicable from this date.

2025.11.12 Legislative Notice for Enforcement Decree (Presidential Decree No. 36053)

Underwent a 40-day legislative notice period (until Dec 22).
Opinions were collected over 70 times across industry, academia, and civil society. Practical details such as high-impact AI determination criteria, transparency obligation implementation methods, and domestic agent designation criteria began to be finalized.

2026.1.20 Amended AI Basic Act Promulgated (Act No. 21311)

Reorganized the National AI Committee into the National AI Strategy Committee, and added provisions promoting public sector AI adoption, establishing legal grounds for AI research institutes, and guaranteeing accessibility for vulnerable groups. Opened new opportunities for business operators seeking to enter the B2G (business-to-government) market.

2026.1.21 Enforcement Decree Promulgated (Presidential Decree No. 36053)

Comprises 33 main articles. Finalized the 11 high-impact AI domains, high-performance AI compute threshold (10^{26} FLOPs or more), domestic agent designation requirements (meeting one of: global revenue of 1 trillion KRW or more, domestic AI revenue of 10 billion KRW or more, or daily average domestic users of 1 million or

more), and transparency implementation methods.

2026.1.22 Phase 1 Enforcement of AI Basic Act

Prior notification obligation, output labeling obligation, high-impact AI risk management, high-performance AI safety assurance, and overseas business operators' domestic agent designation obligation take effect. Distributed the 5 major guidelines (transparency, safety, high-impact judgment, business operator duties, impact assessment) simultaneously. Opened the AI Basic Act Support Desk. A guidance period of at least one year is granted prior to the imposition of administrative fines.

2026.5.14 Colorado, US Enacts Replacement AI Legislation (SB 26-189)

The Colorado AI Consumer Protection Act (SB 24-205) enacted in 2024 was originally scheduled to take effect on 2026.2.1, but its effective date was postponed to 2026.6.30 by special session legislation in August 2025, and prior to enforcement on this day, it was replaced by a new law focused on disclosures and consumer rights regarding Automated Decision-Making Technology (ADMT). The new law takes effect on 2027.1.1. Korean companies providing

services in the State of Colorado must prepare according to the new law.

2026.7.20 Amended Enforcement Decree Promulgated (Takes Effect the Following Day)

Included a confirmation system (Article 15, Paragraphs 4 and 5 of the Enforcement Decree) designating AI products and services confirmed by the Minister of Science and ICT as candidates for priority consideration in public procurement, and added history of administrative fines (Article 29, Paragraph 1, Subparagraph 4 of the Enforcement Decree) to domestic agent designation criteria, bringing the total to four criteria.

2026.7.21 Amended AI Basic Act Takes Effect

Obligation of state agencies, etc., to give priority consideration to confirmed AI products and services (Article 16, Paragraph 3) and immunity clause for procurement officials (Article 16, Paragraph 4) take effect. Procedures for establishing AI research institutes and legal grounds for financial support for vulnerable groups are detailed. Scope of vulnerable groups expands to include career-interrupted women, job seekers, employees of non-capital-region SMEs, and

farmers/fishers.

2026.8.2 EU AI Act General Application Begins: Chatbot Disclosure, Deepfake Labeling, GPAI Enforcement Authority

Transparency obligations (Article 50) apply, requiring AI chatbots to disclose to users that they are interacting with AI, and attaching labeling obligations to deepfake content. Enforcement authority and sanctions regarding general-purpose AI also become operational on this date. Korean business operators providing services to the EU market are also subject to these requirements.

2027.12.2 EU AI Act Standalone High-Risk AI (Annex III) Obligations Take Effect (Omnibus Postponement)

Originally scheduled to take effect on 2026.8.2, but postponed to 2027.12.2 due to the Digital Omnibus amendment politically agreed upon by the Council and European Parliament in May 2026.

Conformity assessment, risk management, and documentation obligations are imposed on standalone high-risk AI systems in domains such as biometrics, recruitment, loan screening, educational evaluation, and law enforcement. In case of violation, administrative fines up to 35 million EUR or 7% of global annual turnover.

2028.8.2 EU AI Act Embedded High-Risk AI (Annex I) Obligations Take Effect (Omnibus Postponement)

Third-party conformity assessment obligations are imposed when AI is embedded in products subject to EU product safety legislation, such as medical devices, machinery, toys, and aviation equipment. Originally scheduled for enforcement on 2027.8.2, it was postponed by the same Omnibus amendment. Korean manufacturers exporting such products to the EU must establish a conformity assessment framework by this point.

Appendix E. Sample Document Templates for Risk Management Measures and Impact Assessments

This appendix provides sample templates for the three core documents that high-impact AI business operators must prepare. Examples have been populated using a hypothetical AI credit scoring system (hereinafter "CreditScore AI") as a case study. In practice, these may be modified and used according to your system.

Template 1. Risk Management Measures (Article 34, Paragraph 1, Subparagraph 1 of the AI Basic Act)

A. Basic Information

Business Operator Name: Fintech Solution Co., Ltd.

AI System Name: CreditScore AI v2.3

Classification: High-Impact Artificial Intelligence (Article 2, Subparagraph 4, Item g: Decisions or evaluations that significantly affect individual rights and obligations, such as employment and loan screening)

Date of Preparation: August 1, 2026

Prepared by: OO Kim, Team Leader, AI Ethics Team

Final Approver: OO Park, CEO

Document Version: 1.0

Next Scheduled Update Date: February 1, 2027 (6-month cycle)

B. System Overview

(1) Function Description

CreditScore AI is a system that receives personal credit-related data as input and recommends loan approval status and interest rate

levels. Bank loan screening officers use it as reference material.

(2) Purpose of Use

Calculates credit scores, repayment capabilities, and default probabilities during personal credit loan screening processes to propose loan approval decisions and appropriate interest rate bands.

(3) Target Users

Loan screening personnel at 12 partner financial institutions
(approx. 340 persons)

(4) Expected Utilization Scale

Average monthly loan screenings: approx. 15,000 cases. Utilized as reference material for approximately 180,000 decision-making cases annually.

(5) Training Data Outline

Approximately 3.2 million de-identified loan history records accumulated from 2018 to 2025. Merged statistical data provided by the Korea Credit Information Services. Personal information

processing is based on Article 32 of the Use and Protection of Credit Information Act and Article 28-2 of the Personal Information Protection Act (Special Provisions on Pseudonymized Information).

C. Risk Identification

(1) Discrimination Risk

Identified Risk: Potential for loan approval rate disparities based on gender, age, or region of residence

Severity: High

Likelihood: Medium

Grounds: Past lending practice biases may be reflected in the training data. A statistically significant higher loan rejection pattern for residents of certain regions was discovered during a 2025 internal audit.

(2) Privacy Risk

Identified Risk: Potential for sensitive information (medical expense spending patterns, religious organization donations, etc.) to be

indirectly reflected during the credit score calculation process

Severity: Medium

Likelihood: Low

Grounds: Although sensitive information is not directly used in input variables, consumption pattern data may act as proxy variables.

(3) Safety Risk

Identified Risk: Financial loss caused by unqualified loan approvals in the event of system error

Severity: Medium

Likelihood: Low

Grounds: System outputs are reference materials and final decisions are made by officers. However, the rate of officers following system recommendations as-is is high at 87 percent, effectively approaching automated decision-making.

(4) Explainability Risk

Identified Risk: Situations where users cannot understand the reasons for rejection upon loan denial

Severity: High

Likelihood: High

Grounds: The current model is XGBoost-based and uses non-linear combinations of hundreds of variables, making it difficult to explain the contribution of individual variables to users.

D. Risk Mitigation Measures

(1) Technical Measures

Fairness Testing: Measure approval rate disparities by gender/age/region once per quarter. If disparity exceeds 5 percentage points, retrain model or apply adjustment weights.

Explainability Tools: Calculate SHAP (SHapley Additive exPlanations) values to present the top 3 influential factors for rejected cases.

Anomaly Detection: Monitor approval rates, average interest rates, and rejection reason distributions on a daily basis. Trigger automated alert if deviating by 10 percent or more compared to the previous month.

(2) Organizational Measures

The AI Ethics Team (3 members) prepares quarterly reports and submits them to the Compliance Officer.

External Audit: Request fairness and accuracy verification from an independent third-party institution once a year.

Internal Training: Conduct training on AI system limitations twice a year for loan screening officers.

(3) Monitoring Plan

Real-Time Dashboard: Display approval rate, rejection rate, average processing time, and number of objections in real time.

Incident Reporting System: Separately categorize AI-related user complaints to analyze monthly trends.

Model Performance Oversight: Measure AUC (Area Under Curve) quarterly. Trigger retraining if dropping below 0.75.

E. Human Intervention System

(1) Types of Decision-Making Requiring Human Intervention

All loan rejection cases, cases applying the highest interest rate band (15 percent per annum or higher), and received objection cases

(2) Intervention Method

Loan screening officer reviews AI recommendation results and SHAP explanations before making final approval/rejection decisions. If overriding AI recommendations, record override reasons in the system.

(3) Designated Personnel

Positions of Team Leader of Loan Screening or higher at each partner financial institution. Override authority is restricted to team leader level.

(4) Objection Procedure

Users notified of loan rejection may file an objection through the relevant financial institution's customer service center within 14 days. Objection cases bypass the AI system and are directly re-screened by a human underwriter.

F. Explanation Measures

(1) Scope of Explanation to Users

Upon loan rejection: Provide the top 3 factors influencing rejection in natural language. E.g., "Number of new loans in the past 6 months is 3 or more", "Ratio of existing loan repayments to annual income exceeds 40 percent".

Upon loan approval: Provide the top 2 major factors influencing the calculation of the applied interest rate.

(2) Method of Explanation

Written notification (mail or email) + detailed view screen within the mobile app.

(3) Notice on Explanation Limitations

State at the bottom of the explanation document that SHAP value-based explanations are approximations and do not fully reflect interactions between individual variables.

G. Document Management

Retention Period: 5 years (statutory retention period under Article 27, Paragraph 2 of the Enforcement Decree, including electronic storage)

Update Cycle: 6 months (ad-hoc updates upon model updates, statutory changes, or incident occurrences)

Department in Charge: AI Ethics Team

Storage Method: Stored with version control in the internal Document Management System (DMS). Accessible during external audits.

Template 2. Impact Assessment Report (Article 35 of the AI Basic Act)

A. Assessment Overview

Target of Assessment: CreditScore AI v2.3

Assessment Period: June 1, 2026 – July 15, 2026 (45 days)

Assessors: AI Ethics Team + 2 external advisory members (1 law professor, 1 data science professor)

Purpose of Assessment: Conduct an autonomous impact assessment pursuant to Article 35 of the AI Basic Act

B. Social Impact Analysis

(1) Positive Impacts

Reduced loan screening processing time from an average of 3 business days to 4 hours. Improved consistency in screening criteria, reducing officer-to-officer variance. Expanded accessibility to micro-loans (5 million KRW or less): Micro-loan approval cases increased by 22 percent after AI adoption.

(2) Negative Impacts

Credit evaluation accuracy for non-standard income earners such as self-employed individuals and freelancers is lower compared to

full-time employees. Potential exclusion of elderly users (aged 65 or older) from AI-based screening benefits due to difficulties in accessing digital channels. Concerns over weakening officers' independent screening capabilities due to deepening system reliance.

(3) Affected Groups

Direct Impact: Approximately 180,000 loan applicants annually

Indirect Impact: Household members and guarantors of loan applicants

C. Fundamental Rights Impact Analysis

(1) Privacy (Right to Personal Information Self-Determination)

Direct identification is not possible due to pseudonymized information processing, but consumption pattern analysis may constitute indirect profiling. A separate Personal Information Impact Assessment (Article 33 of the Personal Information Protection Act) was conducted and passed with a compliant rating (May 2026).

(2) Non-Discrimination (Right to Equality)

As a result of fairness testing, the approval rate disparity between genders was 2.1 percentage points, and between age groups was 3.8 percentage points, within the threshold (5 percentage points). However, the approval rate for residents of certain regions (rural areas) was 7.2 percentage points lower than for urban residents, so corrective measures are being implemented.

(3) Right to Receive Explanation

Top 3 rejection reason factors are provided through the introduction of a SHAP-based explanation system. However, in a user comprehension survey, only 43 percent responded that they "fully understood the explanation." Work on improving explanation documents is underway.

(4) Fair Procedure (Due Process)

Establishment of objection procedure completed. Out of 47 objection cases received in the first half of 2026, original decisions were overridden in 12 cases (25.5 percent).

D. Environmental Impact Analysis

Power consumption per training run: approx. 320 kWh (8 GPU servers, 72 hours)

Annual inference power consumption: approx. 1,200 kWh

Carbon emission conversion: approx. 0.54 metric tons CO2 annually (based on Korea's power grid carbon emission factor)

Comparison: Significantly lower energy consumption compared to human workforce performing the same tasks, but cumulative environmental load needs consideration if model retraining frequency increases.

E. Overall Assessment and Improvement Plan

Overall Rating: Conditionally Compliant

3 Major Improvement Tasks:

Task 1. Correction of approval rate disparity for rural residents (Target: within 5 percentage points, Deadline: December 2026)

Task 2. Improvement of user explanation comprehension (Target: 60 percent or higher "fully understood", Deadline: March 2027)

Task 3. Advancement of evaluation model for non-standard income earners (Target: accuracy AUC 0.80 or higher for self-employed, Deadline: June 2027)

F. Stakeholder Consultation Results

(1) Internal Stakeholders

Survey of loan screening officers (127 respondents): 78 percent "satisfied with accuracy of AI recommendations", 34 percent "have overridden AI recommendations", 21 percent "feel burdened when overriding".

(2) External Stakeholders

Consumer group roundtable (June 20, 2026, 3 participating groups): Pointed out lack of specificity in explanations, requested strengthening promotion of objection procedures.

Financial Supervisory Service prior consultation (July 3, 2026): Recommended shortening fairness testing cycle from quarterly to monthly.

(3) Reflection Results

Consumer group feedback: Reflected in explanation document improvement work (Improvement Task 2).

Financial Supervisory Service recommendation: Scheduled to transition to monthly fairness testing starting in Q4 2026.

Template 3. Safety Check Log (Article 34, Paragraph 1, Subparagraphs 1 and 5 of the AI Basic Act)

A. Log Format

Date/Time | Check Item | Check Result | Action Taken | Person in Charge

2026.08.01 09:00 | Model inference accuracy (AUC) | 0.82 | Normal range (threshold: 0.75 or higher) | OO Kim

2026.08.01 09:00 | Daily approval rate | 67.3% | Normal range compared to previous month (65.8%) | OO Kim

2026.08.01 09:00 | Average response time | 1.2 sec | Within threshold (3 sec) | OO Kim

2026.08.01 09:00 | Anomaly transaction detection | 0 cases | Not applicable | OO Kim

2026.08.05 14:30 | Gender approval rate disparity | 2.3%p | Within threshold (5%p) | OO Park

2026.08.05 14:30 | Age group approval rate disparity | 4.1%p | Within threshold (5%p), entered caution band | OO Park

2026.08.05 14:30 | Regional approval rate disparity | 6.8%p | Exceeded threshold, applied adjustment weight action | OO Park

2026.08.12 10:00 | System availability | 99.7% | Met threshold (99.5%) | OO Lee

2026.08.12 10:00 | Number of objection filings | 8 cases | Slight increase compared to previous month (6 cases), cause analysis | OO Lee

2026.08.15 16:00 | Emergency check: Inference delay occurred | Response 8.7 sec | GPU server Node #2 overload, traffic distributed | OO Kim

2026.08.15 17:30 | Emergency check: Confirmation after action |
Response 1.4 sec | Normal recovery confirmed | OO Kim

2026.08.19 09:00 | Model drift inspection | KL 0.03 | Within threshold
(0.1), normal | OO Park

2026.08.26 09:00 | Monthly comprehensive check | Compliant |
August comprehensive report preparation completed | OO Kim

B. Inspection Frequency Standards

Daily Inspection: Inference accuracy, approval rate, response time,
anomaly transactions

Weekly Inspection: Fairness metrics (gender, age, region), objection
status

Monthly Inspection: Model drift, system availability, comprehensive
report

Ad-hoc Inspection: System failures, surges in user complaints,
external environmental changes (statutory amendments, rapid
changes in economic indicators)

C. Retention and Management

Retention Period: 5 years (statutory retention period under Article 27, Paragraph 2 of the Enforcement Decree, including electronic storage)

Retention Format: Electronic log files (CSV or JSON) + monthly comprehensive report (PDF)

Access Rights: AI Ethics Team, Compliance Officer, External Auditor

Remarks: Emergency check cases are verbally reported to the AI Ethics Team Leader within 24 hours of occurrence, and a written report is prepared within 48 hours.

D. Incident Reporting Template (In Case of Major Incidents)

Date and Time of Incident:

Circumstances of Discovery:

Scope of Impact (Estimated number of affected cases):

Emergency Action Taken:

Cause Analysis (Initial):

Recurrence Prevention Measures:

Reporter:

Date and Time of Report:

Example Entry:

Date and Time of Incident: September 3, 2026, 11:20

Circumstances of Discovery: Abnormal pattern detected on monitoring dashboard showing approval rate sharply dropping by 30 percentage points compared to the previous day

Scope of Impact: Affected approx. 120 cases processed during that time window

Emergency Action Taken: 11:35 Temporarily suspended AI system recommendation feature, switched to manual screening. 12:00 Initiated cause investigation.

Cause Analysis (Initial): Variable mapping error occurred during previous night's batch update. Column order of income variable and

debt variable was swapped.

Recurrence Prevention Measures: Introduce automated regression testing after batch updates. Add variable mapping checksum verification procedures.

Reporter: OO Kim, AI Ethics Team

Date and Time of Report: September 3, 2026, 15:00

Copyright Page

AI Basic Act Business Operator Practical Commentary

ISBN |

Author | Gyeong-Jin Kim

Publisher (Individual) | Gyeong-Jin Kim

Publisher (Entity) | Attorney Gyeong-Jin Kim Publishing House

Publisher Registration | March 10, 2025 (No. 2025-000015)

Address | Room 304, Baek-il Building, 91 Jeonnong-ro,
Dongdaemun-gu, Seoul

Phone | +82-2-6338-1905

Email | kimkj008@gmail.com

Price: 25,000 KRW

© Gyeong-Jin Kim 2026

This book is the intellectual property of the author and
unauthorized reproduction or copying is prohibited.

Note) AI tools were used to assist in part of the content and editing
process of this book.

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.



PayPal

[**https://paypal.me/kimkjj**](https://paypal.me/kimkjj)

Scan the QR code or open the link above.