

The EU's General Purpose Artificial Intelligence Code of Practice

Report prepared for the National Institute of
Informatics - Japan

Rapporteur

David Restrepo Amariles, HEC Paris

Contributor

Damien Charlotin, HEC Paris

Information on this Report

This report was prepared at the request of the National Institute of Informatics of Japan. David Restrepo Amariles, Professor at HEC Paris, served as Rapporteur and led the design and preparation of the report. Damien Charlotin, Researcher and Co-Director of the Smart Law Hub at HEC Paris, served as Contributor, providing substantive research and drafting support. Both are to be regarded as authors of this report. The report is intended to serve as a resource for understanding the European Union's General-Purpose AI Code of Practice and to provide analysis of direct relevance to Japanese firms.

Disclaimer: This report does not constitute legal advice and should not be relied upon as an official interpretation of the European Union's General-Purpose AI Code of Practice. The views and analysis expressed in this document do not represent the position of any institution or body of the European Union, nor does this report constitute an official document of the European Union.

Recommended citation: *Restrepo Amariles, D., & Charlotin, D. The EU's General-Purpose Artificial Intelligence Code of Practice: Report to the National Institute of Informatics of Japan (2026).*

Executive Summary

The EU's AI Act, which entered into force in August 2024, represents the world's first comprehensive horizontal regulation of artificial intelligence. For Japanese companies active in the EU market – whether as developers of AI models, manufacturers of AI-enabled products, or deployers of third-party AI systems – it creates a new and demanding compliance landscape that is already in effect and will reach full force by August 2027.

This briefing focuses on the regulatory framework governing general-purpose AI (“GPAI”) models: large-scale foundation models capable of performing a wide range of tasks. The AI Act regulates GPAI at the model layer rather than the use-case layer, meaning obligations attach to the provider of a model regardless of its ultimate application. Two instruments operationalise these obligations: the GPAI Code of Practice (July 2025), a voluntary compliance tool offering a presumption of conformity with the Act's transparency, copyright, and safety requirements; and the Commission Guidelines (July 2025), which clarify the scope of GPAI duties.

The Code has attracted 28 signatories to date, including all major US technology companies except Meta. Non-signatories remain subject to the same underlying legal obligations but face greater scrutiny from the European AI Office. For Japanese companies, the decision of whether to sign – or to demonstrate compliance through alternative means – is one of the most immediate strategic choices the framework presents.

This briefing closes with seven strategic recommendations for Japanese companies and a compliance blueprint addressing the four principal challenges ahead: mapping AI system footprints, navigating the high-risk conformity assessment process, building AI governance infrastructure, and achieving audit readiness for engagement with EU supervisory authorities. The complexity and ongoing evolution of the framework make specialist legal and technical guidance an essential investment for any organisation seeking to compete in the EU AI market with confidence.

Table of Contents

INFORMATION ON THIS REPORT	2
EXECUTIVE SUMMARY	3
TABLE OF CONTENTS	4
BACKGROUND AND GENESIS	6
A. The AI Act	6
i. The AI Act's legislative journey	6
ii. Towards a Code of Practice	7
B. Drafting the Code of Practice	8
C. Signatories and non-signatories	10
D. The EU Commission Guidelines for GPAI Providers	11
THE GPAI CODE OF PRACTICE	14
A. Main Principles	14
B. Timelines	15
C. In details	16
i. Transparency	17
ii. Copyright	20
iii. Safety and security	21

D. Future Outlooks	24
KEY CONSIDERATIONS FOR THE APPLICATION OF THE CODE OF CONDUCT	26
A. The value chain	26
i. Downstream uses	26
ii. The Risks of Model Modification and Fine-Tuning	27
B. Territorial Scope: The Brussels Effect	27
C. A clash of philosophies ?	29
D. Sectoral Impacts: Robotics	31
STRATEGIC GUIDANCE FOR JAPANESE COMPANIES	32
1. Establish a Dual-Track Governance Strategy	32
2. Implement a Global Copyright Opt-Out Protocol	32
3. Rigorous Monitoring of Compute Thresholds	33
4. Leverage the JETRO and Digital Partnership Networks	33
5. Proactive Engagement in the Standard-Setting Process	33
6. Strengthen Contractual Safeguards in the AI Supply Chain	34
7. Proactive Engagement in the Standard-Setting Process	34
8. Adopt the right mindset	35
FURTHER GUIDANCE TO COMPLY WITH THE AI ACT	37
Challenge 1: Mapping your AI footprint	37
Challenge 2: Navigating the high-risk regime	38
Challenge 3: Building a governance infrastructure	38
Challenge 4: Audit readiness and engagement with supervisory authorities	38

Background and Genesis

A. The AI Act

The origins of the GPAI cannot be understood without first looking at the process of drafting and enacting the AI Act.¹

i. The AI Act's legislative journey

The European Commission published its proposal for an Artificial Intelligence Act on 21 April 2021. The original framework envisioned four tiers of regulatory intensity – from prohibited practices at the top to minimal-risk systems largely left unregulated – with obligations calibrated to each system's potential impact on health, safety, and fundamental rights.

What the Commission's original draft did not anticipate, however, was the seismic shift in the AI landscape that would occur eighteen months later: the public release of ChatGPT in November 2022, which represented a significant breakthrough in generative AI.² Multiple “chatbots” soon entered the market, becoming the technology with the fastest rhythm of adoption ever.

For the EU, this proved an issue: the AI Act's fundamental risk-based framework was designed around AI systems with specific, identifiable use cases. Yet, foundation models are trained on vast quantities of data to perform a wide range of tasks and

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, 13 June 2024 (“**AI Act**”), available [here](#).

² Berkman Klein Center, “The EU AI Act – A Real-Time Experiment to Regulate Generative AI”, 16 November 2023, available [here](#).

can be integrated into countless downstream applications. They did not fit neatly into the Commission’s use-case-specific risk categories.

While the legislative process kicked in and sought to plug the gap, the regulation of foundation models became one of the most controversial aspects during the trilogue negotiations – with Germany, France, and Italy eventually lobbying for an approach relying more on “self-regulation” and less on hard-coded rules.³

The final text uses the notion of “general-purpose AI”, defined as AI that has been trained with a large amount of data, is capable of performing a wide range of different tasks, and can be integrated into a variety of downstream systems or applications.⁴

Per Chapter V of the final text, which governs such models, a tiered system is applicable: all GPAI providers must meet baseline transparency and documentation requirements, while providers of models with “systemic risk” – determined primarily by a computational threshold of 10^{25} floating-point operations (“**FLOPs**”) – face additional obligations including model evaluations, adversarial testing, and incident reporting.

ii. Towards a Code of Practice

From the experience of drafting and enacting the AI Act, the EU legislator recognized that this regulatory area could not keep easily pace with technical shifts in model architecture or training methodologies. This is what prompted the EU legislator to introduce the Code as a dynamic living document that translates the high-level obligations of Articles 53 and 55 into concrete, actionable measures. As

³ TechPolicy.press, “Will Disagreement Over Foundation Models Put the EU AI Act at Risk?”, 29 November 2023, available [here](#).

⁴ AI Act, Article 2(63).

well, the AI Act was significantly clarified by the EU’s subsequent Guidelines in this respect (“**Guidelines**”).⁵

In the AI Act’s architecture, GPAI is regulated at the *model layer* rather than the *use-case layer*. This is structurally important for compliance strategy, as it means that the GPAI duties attach to the provider of a model “regardless of what is or will be its ultimate use,”⁶ while separate requirements attach to AI systems (and their risk category) depending on deployment context.

B. Drafting the Code of Practice

The development of the GPAI Code of Practice followed an unprecedented multi-stakeholder process facilitated by the European AI Office, which was established within the European Commission in May 2024 to centralize oversight of GPAI models.

⁵ Guidelines on the scope of the obligations for general-purpose AI models established by Regulation (EU) 2024/1689 (AI Act), C(2025) 5045 final, 18 July 2025.

⁶ European Commission, General-Purpose AI Models in the AI Act – Questions & Answers, 9 September 2025, available [here](#).

Timeline

This is the overall drafting process that has taken place until the publication of the Code.



Figure 1: Drafting Timeline (Source: EC)

On 30 September 2024, the AI Office hosted the kick-off event for the Code of Practice Plenary, involving nearly 1,000 participants, including a high number of professional organisations.

As the main addressees of the Code, providers of general-purpose AI models were invited to dedicated workshops with the Chairs and Vice-Chairs, contributing to informing each iterative drafting round, in addition to their Plenary participation. In addition, the

Chairs invited civil society organisations

and downstream industry to additional workshops to allow for even more targeted interactions. The AI Office also invited other public bodies and agencies from all over the world working on risk assessment and mitigation for general-purpose AI models to the Plenary as observers. The drafting process ultimately involved more than 1,400 stakeholders from industry, civil society, and academia.⁷

⁷ See European Commission, Questions and answers on the code of practice for General-Purpose AI, 11 July 2025, available [here](#).

The Draft went through four major iterations between 2024 and 2025. This process was intended to ensure that the requirements for transparency, copyright compliance, and safety were not only legally robust but also technically feasible for industry actors.

C. Signatories and non-signatories

At the time of writing this report, 28 organisations have signed the Code of Practice.⁸ The current tally includes major technology companies such as Amazon, Anthropic, Google, IBM, Microsoft, and OpenAI.

Notably, European AI champions were among the first to commit. Prominent European AI companies that have signed include Aleph Alpha and Mistral AI. This is significant given that these same companies had been at the centre of the lobbying effort by France, Germany, and Italy during the negotiations to water down the AI Act's GPAI provisions. Their subsequent embrace of the Code suggests an acceptance that engagement with the regulatory framework is preferable to resistance.

Not everyone has opted to subscribe in full to the Code: xAI agreed to sign only the chapter covering safety and security, meaning that it will have to demonstrate compliance with the AI Act's obligations concerning transparency and copyright via alternative adequate means. In a post on X (formerly Twitter), Elon Musk opined that “[w]hile the AI Act and the Code have a portion that promotes AI safety, its other parts contain requirements that are profoundly detrimental to innovation and its copyright provisions are clearly an over-reach.”

As of January 2026, Meta is the most notable company that has declined to sign the Code. Major Chinese AI companies, such as Alibaba, Baidu, or DeepSeek, have also

⁸

European Commission, The General-Purpose AI Code of Practice, last updated 2 February 2026, available [here](#).

not signed. The absence of Chinese providers is unsurprising given the geopolitical context, but it raises questions about the Code's effectiveness as a truly global standard-setting instrument.

The divergent approaches of major international players present both risks and opportunities for Japanese companies. The fact that the overwhelming majority of leading GPAI providers – including all major US technology companies except Meta – have signed the Code suggests that it is becoming the *de facto* global standard for responsible AI development. For Japanese firms seeking to partner with or integrate models from these providers, this alignment provides a degree of assurance regarding upstream compliance. Yet, the stance of non-signatories like Meta also demonstrates that alternative compliance pathways remain available under the AI Act. Japanese companies developing their own GPAI models will need to make a strategic decision: sign the Code and benefit from the presumption of conformity and reduced regulatory burden, or maintain flexibility by demonstrating compliance through alternative means – at the cost of increased scrutiny from the AI Office.

D. The EU Commission Guidelines for GPAI Providers

Alongside the Code of Practice, the European Commission published in July 2025 the Guidelines on the scope of the obligations for general-purpose AI models, which serve a different but complementary function: they clarify the Commission's interpretation of who is subject to the GPAI obligations in the first place, and what those obligations require. Although not legally binding the Guidelines set out the Commission's position on which it will base its enforcement action, making them a crucial reference document for any GPAI provider active in the EU market.

The Guidelines introduce a compute-based indicative criterion for determining whether a model qualifies as a GPAI model: training compute greater than 10^{23} FLOP, combined with the ability to generate language (text or audio/speech), text-to-image, or text-to-video outputs. This threshold corresponds to the approximate compute used to train a model with one billion parameters on a large dataset. Crucially, this criterion is indicative rather than determinative. This matters for Japanese companies that develop smaller, domain-specialised models: the compute test provides only a relative safe harbour.

The Guidelines clarify the procedure by which a GPAI model is classified as having “systemic risk” – a status that triggers the most demanding obligations under the AI Act. A model is presumed to have systemic risk when its cumulative training compute exceeds 10^{25} FLOP. Providers who foresee reaching this threshold must notify the Commission without delay, and in any event within two weeks. Importantly, the Guidelines clarify that this notification obligation may arise *before* training is complete, if the provider can reasonably foresee the threshold will be met. Providers may contest classification by submitting substantiated arguments that their model does not, despite its compute level, present systemic risks. Notably, mitigations already put in place are not valid grounds for contesting classification: a model that poses systemic risks must continue to be assessed and mitigated regardless of safety measures implemented.

The Guidelines further address two issues of particular relevance to Japanese firms operating within global AI supply chains. First, on downstream modifiers – entities that fine-tune or otherwise modify a pre-existing GPAI model – the Commission provides an indicative threshold: a modifier becomes a “provider” subject to GPAI obligations only if the compute used for the modification exceeds one third of the

compute used to train the original model. Second, the Guidelines detail the conditions for the open-source exemption from certain transparency obligations. Critically, the exemption does not cover copyright obligations: even open-source GPAI providers must maintain a copyright compliance policy and publish a training data summary.

The GPAI Code of Practice

The General-Purpose AI Code of Practice is the AI Act’s main “soft-law” bridge between legally binding Articles 53 and 55 and operational compliance. Although the Commission describes it as a “voluntary tool”, it comes in support – and fleshes out – actual obligations from the Act.

A. Main Principles

The Code has three chapters. The first two, Transparency and Copyright, apply to all GPAI model providers. The third, Safety and Security chapter, only applies to providers of GPAI models with systemic risk (above the 10^{25} FLOP threshold), currently a small group of 5-15 companies worldwide. The Code needs to be read and understood in light of the accompanying Guidelines.

Adherence to the Code is voluntary and provides a rebuttable presumption of conformity. This creates a strong incentive for model providers to sign on, as EU regulators are likely to treat adherence to the Code as evidence of compliance with the AI Act.

The practical implications of this design are significant:

STATUS	REGULATORY TREATMENT	ADMINISTRATIVE BURDEN
Signatory	Presumption of conformity; Commission focuses enforcement on monitoring adherence	Reduced; standardised documentation via Model Documentation Form
Non-Signatory	Must independently demonstrate compliance through alternative means	Increased; subject to more information requests and heightened scrutiny

Table 1: Practical Implications

The Code’s voluntary nature also serves important political economy functions. By allowing providers to demonstrate compliance through alternative means, the Commission preserved flexibility for companies with legitimate concerns about specific provisions – such as xAI and Meta’s objections to the copyright chapter – while maintaining the integrity of the underlying legal obligations. While adherence to the Code is voluntary, compliance with the AI Act is not.

B. Timelines

The following timelines govern the phased implementation of GPAI obligations:

IMPLEMENTATION PHASE	DATE	REGULATORY IMPLICATION
Entry into Force	August 1, 2024	The AI Act officially becomes part of EU law.
Initial GPAI Obligations	August 2, 2025	Rules on transparency and copyright apply to new models.
AI Office Enforcement Powers	August 2, 2026	The Commission can begin imposing fines for GPAI violations.
Full Application of High-Risk Rules	August 2, 2026	Comprehensive requirements for high-risk systems enter into force.
Legacy GPAI Deadline	August 2, 2027	Existing models (placed before Aug 2025) must achieve full compliance.

Table 2: Implementation timeline

The GPAI rules took effect on 2 August 2025, meaning all new models released from that date must comply. However, the Commission’s enforcement actions – such as

requests for information, access to models, or model recalls – will only begin a year later. This grace period gives providers time to work with the AI Office to ensure they meet the standards.

As clarified by the Commission in the Q&A related to the Code, if providers “do not fully implement all commitments immediately after signing the Code, the AI Office will not consider them to have broken their commitments under the Code and will not reproach them for violating the AI Act. Instead, in such cases, the AI Office will consider them to act in good faith and will be ready to collaborate to find ways to ensure full compliance.”⁹

For Japanese firms, the August 2026 milestone is particularly critical, as it marks the point at which the European Commission’s enforcement powers, including the ability to conduct model audits and issue recalls, come into full effect. Companies that have not yet established robust documentation and risk management frameworks face material regulatory risks, including fines that can reach up to 3% of global annual turnover for GPAI-related infractions or up to 7% for prohibited practices.

C. In details

The Code of Practice translates the high-level obligations of Articles 53 and 55 of the AI Act into concrete, operational requirements.

The three chapters address distinct but interconnected concerns. Transparency requirements (i) ensure that regulators and downstream users have sufficient

⁹ See European Commission, Questions and answers on the code of practice for General-Purpose AI, 11 July 2025, available [here](#).

information to assess a model's risks and capabilities. Copyright provisions (ii) operationalise the EU's commitment to protecting intellectual property rights in the age of generative AI. Safety and security obligations (iii) – applicable only to the most advanced models – establish a comprehensive risk governance framework designed to prevent systemic harms before they materialise.

i. Transparency

Transparency is the fundamental pillar upon which the GPAI regulatory regime is built. The AI Act and its accompanying Code of Practice mandate that providers maintain a high degree of “explainability” and documentation throughout the model's lifecycle. This is designed to solve the “information asymmetry” problem in the AI value chain, where downstream users may not fully understand the risks or limitations of the technology they are deploying.

Under Article 53(1) of the AI Act, all providers of GPAI models must fulfil three primary transparency-related duties:

- the creation of technical documentation for regulators,
- the provision of information to downstream providers, and
- the publication of a summary of the training content.

The technical documentation must be detailed enough for the AI Office and national authorities to assess the model's compliance with safety and fundamental rights standards. It must include the model's architecture, training methodologies, data provenance, and the computational resources used for training – often measured in FLOPs.

a. Training data

A central and mandatory component of the transparency regime is the Training Data Summary (“**TDS**”), for which the European Commission published a standardized template in July 2025, after taking comments from over 100 stakeholders.¹⁰ This template is designed to provide visibility into the data used to train a model without forcing the disclosure of trade secrets or work-by-work listings.

The Explanatory Notice is straightforward in highlighting that the chief beneficiaries of that transparency would be copyrights holder,¹¹ although the Commission also gestured towards other benefits, such as the information available to downstream users or researchers,¹² or enhanced competition.¹³

TDS TEMPLATE SECTION	KEY REQUIREMENTS FOR GPAI PROVIDERS	NARRATIVE AND STRATEGIC FOCUS
I. General Information	Identification of provider, model version, and intended purposes.	Establishing a baseline for traceability and version control.
II. List of Data Sources	Disclosure of public datasets, private/licensed sources, and scraped content.	Transparency regarding the provenance of high-level information.
III. Scraped Data Detail	Names of crawlers used, collection period, and the top 10% of domains scraped.	Demonstrating compliance with data extraction ethics and law.
IV. Data Processing	Descriptions of filtering, cleaning, and respect for opt-out signals.	Ensuring data quality and legal compliance in pre-processing.
V. Specialized Modalities	Handling of image, audio, or synthetic data used in training.	Addressing modality-specific risks like deepfakes or copyright.

Table 3: TDS Template sections

¹⁰ Explanatory Notice and Template for the Public Summary of Training Content for general-purpose AI models, 12 June 2025, available [here](#), para. 5.

¹¹ *Ibid.*, paras. 7-8.

¹² *Ibid.*, paras. 10-11.

¹³ *Ibid.*, para. 12.

The requirement to list the top 10% of domains scraped from the internet (or 1,000 domains for SMEs) is particularly impactful for developers of generative AI models. It would allow rightsholders and regulators to identify if a model has been disproportionately trained on content from specific industries or publishers. Furthermore, the template requires narrative descriptions of how synthetic data was used, which is critical for identifying potential “model collapse” risks or biases introduced by AI-generated training material.

b. Downstream users information

Beyond the public-facing TDS, the Code establishes specific information-sharing obligations toward downstream providers, i.e., the companies and developers who integrate GPAI models into their own AI systems. This “downstream package” is critical for the functioning of the AI value chain, as downstream providers often bear their own compliance obligations under the AI Act (particularly if they deploy the model in high-risk applications) and cannot fulfil these without adequate upstream documentation.

The GPAI Code of Practice formalises the information-sharing process: necessary documentation must be provided to downstream providers within a reasonable timeframe. The Model Documentation Form creates a tiered disclosure system: general information about the model’s architecture and intended uses is shared with downstream providers, while more sensitive detail (such as the precise computation used for training) are accessible only to regulatory authorities.

For Japanese firms operating in the EU market, this creates both obligations and entitlements. A Japanese company using an EU-based foundation model for autonomous driving applications has a right to receive documentation on the model’s accuracy, robustness, and cybersecurity within the 14-day window. Conversely, a

Japanese technology firm providing an LLM API to European startups must be prepared to handle a potentially high volume of technical enquiries and provide standardised documentation that meets AI Act standards.

ii. Copyright

The Copyright chapter of the GPAI Code of Practice is perhaps the most contested and delicate section of the regulatory framework. It seeks to operationalize the mandate of Article 53(1)(c) of the AI Act, which requires providers to put in place a policy to comply with Union copyright law and related rights. For Japanese firms, this is a point of significant friction, as Japan’s domestic copyright laws (specifically Article 30-4 of the Copyright Act) are among the most flexible in the world, allowing for broad use of copyrighted works for machine learning without the need for an opt-out.

In contrast, the EU requires developers to respect the “reservation of rights” expressed by rightsholders. The Code of Practice demands that signatories use state-of-the-art methods to identify and honour these reservations, most notably through the Robot Exclusion Protocol (robots.txt) and machine-readable metadata.¹⁴ Japanese firms placing models on the EU market must, therefore, ensure that their web crawlers and data ingestion pipelines are re-engineered to globally honour these opt-out signals, even if they are not legally required to do so within Japanese territory.

¹⁴ It should be noted that the robots.txt mechanism, originally designed in the 1990s to manage search engine crawling, is a remarkably crude instrument for the sophisticated task of rights management. It operates at the domain level, not the work level; it cannot distinguish between copyrighted and public domain content on the same site; and it offers no way to express conditional permissions (e.g., “training permitted for non-commercial use”). More fundamentally, the opt-out framework inverts the traditional copyright bargain: rather than rightsholders demonstrating infringement, developers must now prove they respected an ever-expanding patchwork of machine-readable signals scattered across the internet.

Beyond the training phase, the Code addresses the potential for a model to generate infringing outputs. Providers must implement technical and organizational measures to prevent “verbatim reproductions” of protected works. In practice, this would likely involve a multi-layered defence strategy including:

1. **Filtering Mechanisms:** Implementing post-processing filters that block the generation of content that matches known copyrighted material.
2. **Prompt Constraints:** Designing the model to refuse prompts that explicitly request the replication of a specific copyrighted work or style.
3. **Exclusion of Infringing Sources:** In line with the Code, make sure to avoid scraping websites that are notorious for copyright infringement (piracy sites).¹⁵

The Code establishes a mandatory “complaint and redress” system, where rightsholders can lodge inquiries if they suspect their content has been misused. Providers are expected to engage in good-faith remediation, which may include investigations into training sets or the implementation of specific blocks.

iii. Safety and security

The most stringent tier of the GPAI Code of Practice is reserved for models that pose “systemic risks”. Under the AI Act, a model is presumed to have systemic risk if its cumulative computing power used for training exceeds 10^{25} FLOPs. These “frontier models” are viewed as having the potential to cause large-scale negative impacts on the internal market, public safety, or democratic processes due to their high capability and widespread adoption.

¹⁵ The European Commission intends to maintain and share lists of such sites to assist developers in maintaining a “clean” training pool.

For Japanese developers of advanced foundation models, the systemic risk obligations represent a significant jump in compliance complexity. Providers must implement a comprehensive risk management framework that identifies, evaluates, and mitigates risks across the entire model lifecycle. This may include:

- **Adversarial Testing and Red-Teaming:** Conducting rigorous internal and independent external testing to identify vulnerabilities, such as the potential for the model to assist in cyberattacks or the creation of biological hazards.
- **Independent Audits:** Unlike standard models, systemic risk models are subject to third-party audits of their safety mechanisms to ensure they meet industry-best standards.
- **Cybersecurity Fortification:** Implementing high-level protection for the model's weights and infrastructure to prevent unauthorized access or modification by hostile actors.

A critical requirement for systemic risk providers is the obligation to notify the AI Office of any “serious incidents” without undue delay. A serious incident is defined as any malfunction or unforeseen behavior that leads to significant harm to health, safety, or fundamental rights. Documentation of these incidents, along with the corrective measures taken, must be maintained and shared with regulators. In late 2025, the Commission published a template report to that end.¹⁶

Furthermore, post-market monitoring is mandatory. Providers must continuously observe the real-world performance of their models to detect emerging hazards that may not have been apparent during pre-deployment testing. For Japanese firms, this

¹⁶ European Commission, “AI Act: Commission publishes a reporting template for serious incidents involving general-purpose AI models with systemic risk”, 4 November 2025, available [here](#).

necessitates the development of sophisticated telemetry systems and user feedback channels that can feed data back into the risk management cycle.

RISK CATEGORY	THRESHOLD / PRESUMPTION	PRIMARY OBLIGATIONS
Non-Systemic GPAI	General-purpose capabilities; <10 ²⁵ FLOPs.	Transparency, TDS Template, Copyright Policy.
Systemic Risk GPAI	> 10 ²⁵ FLOP; High capabilities.	All of the above, plus: Risk Assessment, Red-Teaming, Incident Reporting, Audits.
Downstream Modification	Training volume > 1/3 of 10 ²³ FLOPs.	Legal re-classification as a “Provider” of a new model.

The threshold of 10²⁵ FLOPs is a technical benchmark that Japanese companies must monitor closely. As computing hardware advances and training efficiency improves, more models will naturally cross this line, bringing them into the most regulated tier of the EU market.

Finally, the Code also mandates formal governance arrangements for systemic risk providers. This includes:

- Designated personnel with clear accountability for safety decisions
- Independent internal oversight functions (analogous to risk committees in financial institutions)
- Board-level reporting on systemic risk assessments
- Documentation of decision-making processes for major deployment decisions

For Japanese firms accustomed to consensus-based corporate governance, the requirement for clearly designated individual accountability may require organisational adaptation.

D. Future Outlooks

As of February 2026, the regulatory landscape is in a state of flux. While the GPAI Code of Practice provides an immediate compliance tool, the long-term goal for the European Commission is the development of “harmonized standards” by European standardization organizations like CEN and CENELEC. These standards translate the law into common technical language, offering even greater legal certainty.

However, the development of these standards has faced significant delays.¹⁷ This makes the Code of Practice the primary, and perhaps only, viable route to demonstrating compliance for the next several years. Japanese firms should closely follow the work of Japanese representatives in these standardization committees, as their contributions are essential to ensuring that Japanese engineering excellence is not barred by Euro-centric technical requirements.

In parallel, calls are multiplying for a deliberate watering down of the AI Act, in line with the EU’s newfound simplification agenda.¹⁸ This may explain why, for instance, none of the signatories of the Code of Practice has, so far, published anything in accordance with its strictures. In this respect, the grace periods (up to August 2027 for pre-August 2025 models) have raised questions with respect to the dating of any given model: for instance, the European Commission is currently investigating

¹⁷ Chambers of progress, “On the AI Act, the EU Need Not Move Fast and Break Things”, 28 January 2026, available [here](#).

¹⁸ Financial Times, “EU set to water down landmark AI act after Big Tech pressure”, 7 November 2025, available [here](#).

whether GPT-5, released on August 7, 2025, qualifies as a pre- or post-2025 model for the purpose of the grace periods.¹⁹

The signatories' silence suggests either that compliance is more operationally complex than anticipated, that providers are testing the limits of the Commission's collaborative first-year approach, or that backroom negotiations are underway to soften requirements before public disclosure becomes necessary. For Japanese firms observing from the outside, the lesson is clear: signature and compliance are distinct acts, and the former does not guarantee the latter.

If the current trajectory holds, from August 2, 2026, the EU AI Office will move from a phase of support and guidance to a phase of active enforcement. This includes the power to request access to models, conduct on-site audits, and even order the recall of models that pose an unacceptable risk. For Japanese companies, the “grace period” afforded by the 2025 launch of the Code is a critical window to “stress-test” their internal governance before the threat of fines becomes real.

¹⁹ See, e.g., Risto Uuk, “The EU AI Act Newsletter #86: Concerns Around GPT-5 Compliance”, 15 September 2025, available [here](#).

Key Considerations for the Application of the Code of Conduct

A. The value chain

i. Downstream uses

The GPAI Code of Practice is designed to ensure accountability across the entire AI value chain, from the foundational model developer to the final application deployer. This “cascading responsibility” model is crucial for Japanese companies that often occupy multiple roles in the ecosystem – as both developers of their own models and deployers of third-party tools.

Under the Code, providers are required to establish channels for the timely disclosure of information to downstream users – not as a courtesy, but as a regulatory necessity. Downstream users in the EU are often required to conduct their own conformity assessments for high-risk applications, and they cannot do so without the technical data from the upstream GPAI provider. Signatories to the Code agree to respond to information requests from downstream partners within 14 days, ensuring that the compliance chain remains unbroken.

For instance, in the case of a Japanese automotive company using an EU-based vision model for autonomous driving, this means they have a right to receive documentation on the model’s accuracy, robustness, and cybersecurity. Conversely, if a Japanese tech firm provides an LLM API to European startups, they must be prepared to handle a high volume of technical inquiries and provide standardized documentation that meets AI Act standards.

ii. The Risks of Model Modification and Fine-Tuning

One of the most complex aspects of the GPAI rules involves the modification of existing models, a common practice for firms looking to tailor AI for specific industrial or linguistic needs. If an entity modifies or fine-tunes a GPAI model to such an extent that it changes the model's core capabilities, they may be legally regarded as the "provider" of a new model.

The current Guidelines suggest a compute-based threshold for this re-classification: if the training volume used for modification exceeds one-third of original model's compute, or one-third of the 10^{23} FLOPs threshold to be considered a general-purpose AI model, the entity responsible for the modification assumes the full legal obligations of a GPAI provider. This "compute carry-over" risk necessitates a highly cautious approach to fine-tuning strategies for Japanese SMEs and startups.

B. Territorial Scope: The Brussels Effect

The Code's design reflects broader geopolitical considerations. Given the cost of training a GPAI model with more than 10^{25} FLOPs, it is unlikely providers will develop different models for different jurisdictions just to avoid compliance with the AI Act. This suggests the Code may have *de facto* global effect, as providers implement its requirements across their entire model portfolio rather than maintaining EU-specific variants.

For Japanese firms, this "Brussels Effect" dynamic means that alignment with the Code may become a competitive necessity even for models not directly marketed in Europe, as it increasingly defines the global baseline for responsible AI development. Indeed, jurisdictional application is triggered whenever:

1. An AI system or GPAI model is placed on the EU market.
2. The outputs of an AI system are used within the EU, even if the system is hosted on servers in Tokyo.
3. An AI-enabled service affects individuals or businesses located within the EU.

This creates a significant “compliance gap” for Japanese companies accustomed to a more relaxed regulatory environment at home. While the Japanese government’s AI Guidelines for Business emphasize core principles such as transparency and bias mitigation,²⁰ they lack the specific, prescriptive documentation requirements and the threat of astronomical fines found in the EU Act.

STRATEGIC CHALLENGE	IMPLICATION FOR JAPANESE FIRMS	MITIGATION STRATEGY
Product Redesign	Systems designed for the Japanese market may not meet EU <i>ex ante</i> standards.	Adopt EU compliance as the global R&D baseline.
Information Disclosure	The TDS template requires disclosure of scraped domains and data size.	Develop aggregated reporting methods that protect trade secrets.
Extraterritorial Fines	Penalties up to 3% of global turnover for GPAI violations; up to 7% for prohibited practices.	Implement rigorous internal “unacceptable risk” audits.
Certification Delays	High-risk systems require months of preparation and assessment.	Start conformity assessments at the design phase (Compliance-by-Design).
Language and Documentation	All regulatory filings, TDS templates, and Model Documentation Forms must be in an official EU language (typically English).	Establish bilingual compliance teams; develop translation workflows for technical documentation.

²⁰

METI, AI Guidelines for Business, Ver1.0, 19 April 2024, available [here](#).

Hardware Limitations	FLOP-based thresholds impact compute-intensive frontier research.	Optimize models for efficiency rather than raw compute volume.
----------------------	---	--

Table 4: Strategic Challenges

As such, Japanese firms face several practical hurdles when aligning with the GPAI Code of Practice. The most immediate is the cost of compliance. According to JETRO surveys, securing human resources and managing rising labour costs are already top operational challenges for Japanese-affiliated companies in Europe.²¹ The AI Act adds a new layer of expense, requiring specialized legal counsel, technical auditors, and compliance officers.

Furthermore, the “innovation-first” approach favoured in Japan often relies on rapid iteration and *ex post* feedback. The EU’s requirement for exhaustive technical documentation *before* market entry (*ex ante*) can slow down this development cycle, potentially leading to a “time-to-market” disadvantage compared to competitors operating solely in less regulated jurisdictions.

C. A clash of philosophies ?

The divergence between the Japanese and European regulatory models reflects deeper philosophical differences about the role of the state in technological governance.²² Japan’s approach is fundamentally “enabling,” viewing AI as a national priority to address structural labour shortages and ensure economic security. The Japanese AI Promotion Act functions as a policy framework statute rather than a set of binding prohibitions, focusing on research support and social implementation.

²¹ JETRO, Survey on Business Conditions of Japanese-Affiliated Companies in Europe, available [here](#).

²² See So & Sato Innovative Lawyers, “AI Regulation in the EU and Japan: A Practical Guide for Cross-Border Businesses”, 23 January 2026, available [here](#).

In contrast, the EU approach is “protective,” viewing AI as a potential threat to fundamental rights that must be constrained by rigid legal barriers. While both jurisdictions share a commitment to “human-centric” AI, they utilize different mechanisms for enforcement. Japan relies on reputational mechanisms and administrative guidance, where lapses are addressed through litigation under existing laws rather than AI-specific fines.

Despite these differences, there is a remarkable consensus forming around core safety principles. The G7 Hiroshima AI Process, which Japan took a leading role in developing during its 2023 G7 presidency, serves as a bridge between these regulatory islands. A commonality analysis reveals that approximately 80% of the recommendations in the Hiroshima International Code of Conduct for Organizations Developing Advanced AI Systems overlap in substance with the provisions of the EU AI Act.²³

Both frameworks emphasize risk assessment, red-teaming, cybersecurity, and the need for traceability in training datasets. For Japanese firms, this alignment is positive: by building governance structures that satisfy the Hiroshima Code of Conduct, they may be 80% of the way toward meeting the requirements of the EU GPAI Code of Practice. The remaining 20% – primarily the more prescriptive documentation and the specific EU copyright opt-out mechanisms – can then be addressed, if need be, as regional add-ons.

²³ SaferAI, “G7 Hiroshima AI Process Code of Conduct and EU AI Act GPAI – Commonality Analysis”, October 2025, available [here](#).

D. Sectoral Impacts: Robotics

Japan's stake in the EU regulatory framework for AI extends well beyond software companies and large language model developers. The country is the world's largest producer, and one of the key markets, for industrial robots. National champions – FANUC, Yaskawa, Kawasaki, Denso, and Mitsubishi Electric – together account for a significant share of global industrial robot market revenue; their products are deployed extensively across European manufacturing, and their strategic pivot towards AI-powered robotics is bringing them squarely within the Act's scope.

As these firms integrate large foundation models – whether developed in-house or sourced from third-party GPAI providers – into their robotic systems, they enter the GPAI value chain as downstream users, and potentially as downstream providers if they modify those models significantly for their own uses. This is on top of the AI Act's high-risk regime: robotic systems used in manufacturing, infrastructure, healthcare, and logistics are likely to fall within the AI Act's Annex I and Annex III high-risk categories, meaning they will be subject to mandatory conformity assessments, CE marking, technical documentation requirements, and human oversight mechanisms.

Moreover, with an eye to the future, Japanese companies should start to investigate how these regulatory layers will apply to robotic systems that become increasingly capable of autonomous decision-making in unstructured environments, since the boundary between a robotic AI system and a GPAI model may well become harder to draw.

Strategic Guidance for Japanese companies

Based on the analysis of the GPAI Code of Practice and its interplay with the Japanese regulatory environment, the following best practices are recommended for Japanese firms:

1. Establish a Dual-Track Governance Strategy

Companies should not attempt to create a single, monolithic AI policy. Instead, they should maintain a flexible, principles-based framework for their domestic operations while establishing a separate, prescriptive “EU Compliance Track” for their international exports. This track should include:

- **Detailed Technical Documentation:** Adopting the EU’s template as the internal standard for all models, ensuring that data provenance and training compute are tracked from day one.
- **Compliance-by-Design:** Integrating risk assessment and red-teaming into the earliest stages of the R&D process, rather than treating them as an afterthought.

2. Implement a Global Copyright Opt-Out Protocol

To mitigate the risk of litigation in the European market, Japanese developers should configure their web crawlers to honour robots.txt and machine-readable rights signals globally. Even if a dataset is being used primarily for a model launched in Japan, the eventual “Brussels Effect” means that high-quality, legally clean datasets will be the most valuable and portable assets in the long term.

3. Rigorous Monitoring of Compute Thresholds

R&D teams must be vigilant of the 10^{23} FLOP threshold (which triggers GPAI classification) and the 10^{25} FLOP threshold (which triggers systemic risk obligations). Training strategies should be optimized not just for model performance, but for regulatory category management. If a model is intended for a non-systemic use case, developers should be cautious about crossing the 10^{25} threshold, which brings with it the burden of independent audits and incident reporting.

4. Leverage the JETRO and Digital Partnership Networks

Japanese firms, especially SMEs, should utilize the resources provided by the Japan External Trade Organization and the EU-Japan Centre for Industrial Cooperation. These organizations provide crucial “regulatory intelligence,” helping firms navigate the specific nuances of European national authorities and the latest updates to the Code of Practice.

5. Proactive Engagement in the Standard-Setting Process

The final technical rules of the AI Act will be written in the form of standards. Japanese companies should send their lead engineers to participate in international standard-setting bodies (ISO/IEC) and collaborate with European committees (CEN/CENELEC) where possible. This ensures that Japanese technologies – such as specialized industrial robots or specific automotive safety systems – remain compatible with European legal requirements.

6. Strengthen Contractual Safeguards in the AI Supply Chain

Japanese firms licensing GPAI models from third parties (whether US, European, or Chinese providers) should negotiate robust contractual protections, including:

- **Compliance representations:** Warranties that the upstream model complies with EU AI Act requirements
- **Compute disclosure:** Disclosure of training compute to assess proximity to regulatory thresholds
- **Indemnification:** Protection against regulatory liability arising from upstream non-compliance
- **Audit rights:** The ability to verify compliance claims, particularly for models approaching systemic risk thresholds
- **Update notifications:** Advance notice of model updates that may affect regulatory classification

These provisions are particularly critical for firms that fine-tune or modify third-party models, as modification may trigger “provider” status and associated liability

7. Proactive Engagement in the Standard-Setting Process

The AI Office represents a novel type of EU regulatory body: technically specialist, explicitly committed to a collaborative enforcement posture, and - t least during the current transitional period – genuinely incentivised to support providers in achieving compliance rather than simply penalising non-compliance. The Guidelines make clear that the AI Office’s stated approach is “collaborative, staged, and proportionate,” and the Commission has explicitly stated that it will give particular consideration to providers who proactively engage before formal enforcement begins in August 2026.

For Japanese companies, this window should be treated as a strategic asset rather than a grace period to be waited out.

Even where a company is confident its models fall below the relevant thresholds, establishing a documented record of that assessment and sharing it informally with the AI Office is a form of regulatory capital: it demonstrates good faith, reduces the likelihood of an adverse classification decision later, and positions the company favourably if its compute usage grows over time.

Japanese companies should also consider participating in the AI Office's broader consultative structures. The AI Office actively invites public bodies, industry actors, and international stakeholders to contribute to the ongoing development of codes of practice, harmonised standards, and enforcement guidelines. Participation in these processes offers not only advance sight of regulatory direction, but the opportunity to ensure that Japanese industry's technical realities and compliance constraints are reflected in the rules as they are finalized.

8. Adopt the right mindset

A final strategic consideration: Japanese firms must decide whether to approach EU compliance as a burden to be minimised or an opportunity to be leveraged.

The former mindset leads to grudging, minimal compliance – meeting the letter of requirements while preserving maximum flexibility. The latter mindset recognises that robust AI governance can be a competitive differentiator, signalling trustworthiness to European customers, partners, and regulators. In a market increasingly sensitive to AI risks, “EU-compliant” may become a mark of quality rather than merely a regulatory checkbox. Japanese firms that internalise this perspective – treating the Code of Practice not as an imposition but as a framework

for excellence – may find themselves better positioned than competitors who view compliance as purely a cost centre.

Further Guidance to Comply with the AI Act

The foregoing sections have focused primarily on the GPAI obligations that entered into application in August 2025. These are, however, only one layer of a regulation whose full weight will be felt progressively through to 2027 and beyond. For any Japanese company with a commercial presence in the EU; a structured compliance programme is no longer optional. What follows is a practical blueprint for approaching that programme, organised around the four challenges that, in our assessment, represent the most significant sources of risk and complexity.

Challenge 1: Mapping your AI footprint

The foundational step in any AI Act compliance programme is an inventory of all AI systems and models that fall within the Act's scope. This is harder than it sounds. The Act's definition of an "AI system" is broad – covering any machine-based system that infers outputs such as predictions, recommendations, or decisions from the inputs it receives – and the boundary between a conventional software system and a regulated AI system is not always self-evident.

The mapping exercise must answer three questions for each system: Does it fall within the Act's definition of an AI system ? If so, what risk category does it occupy – prohibited, high-risk, limited-risk, or minimal-risk ? And if it incorporates a GPAI model, who in the supply chain bears the GPAI provider obligations, and has that party demonstrably fulfilled them ? The answers will determine the compliance obligations that apply and the timeline by which they must be met.

Challenge 2: Navigating the high-risk regime

For most industrial companies, the high-risk provisions of the Act will represent the most demanding compliance challenge. High-risk AI systems must satisfy an extensive set of mandatory requirements before they can be placed on the EU market: a documented risk management system, data governance and quality assurance procedures, detailed technical documentation, automatic logging of operations, transparency obligations towards users, human oversight mechanisms, and, where required, a third-party conformity assessment conducted by an accredited notified body before CE marking can be obtained.

Challenge 3: Building a governance infrastructure

Technical compliance with the Act's requirements is inseparable from the organisational infrastructure that sustains it. The Act requires providers and deployers of high-risk AI systems to establish quality management systems covering their AI development and deployment processes – an obligation that closely mirrors the ISO 42001 AI management system standard, which provides a useful framework for companies beginning to build this infrastructure.

The challenge is less one of building from scratch than of extending existing quality and risk frameworks to cover AI-specific obligations, and ensuring that those frameworks are documented in a manner that will satisfy European regulators and notified bodies.

Challenge 4: Audit readiness and engagement with supervisory authorities

The AI Act creates a multi-layered supervisory structure. At the EU level, the AI Office is responsible for GPAI model oversight; national market surveillance

authorities in each Member State are responsible for high-risk AI system oversight; and the European AI Board coordinates between them. From August 2026, these authorities will have broad investigatory and enforcement powers.

Achieving audit readiness requires that documentation is not merely prepared but maintained and updated throughout the AI system's lifecycle. It also requires that the company has designated points of contact for regulatory interactions, and has rehearsed how it would respond to an information request or an on-site inspection. Many Japanese companies operating through European subsidiaries or distributors will need to clarify internally who bears the authorised representative role and ensure that person or entity has the technical competence and legal authority to engage with regulators on the company's behalf.

The AI Act is complex, technically demanding, and still evolving: harmonised standards are not yet finalised, notified bodies are still being designated, and the enforcement practice of national authorities is only beginning to take shape. For companies without in-house EU regulatory expertise, navigating this landscape in real time – while continuing to develop and deploy AI systems commercially – is a significant undertaking.

The four challenges outlined above each require a distinct combination of skills: legal interpretation of a regulation that is simultaneously broad in scope and highly technical in detail; engineering judgment to classify systems correctly and design compliant architectures; familiarity with EU administrative processes and the working methods of the AI Office and national supervisory authorities; and the cross-jurisdictional perspective needed to reconcile EU requirements with domestic Japanese regulatory obligations and international standards.