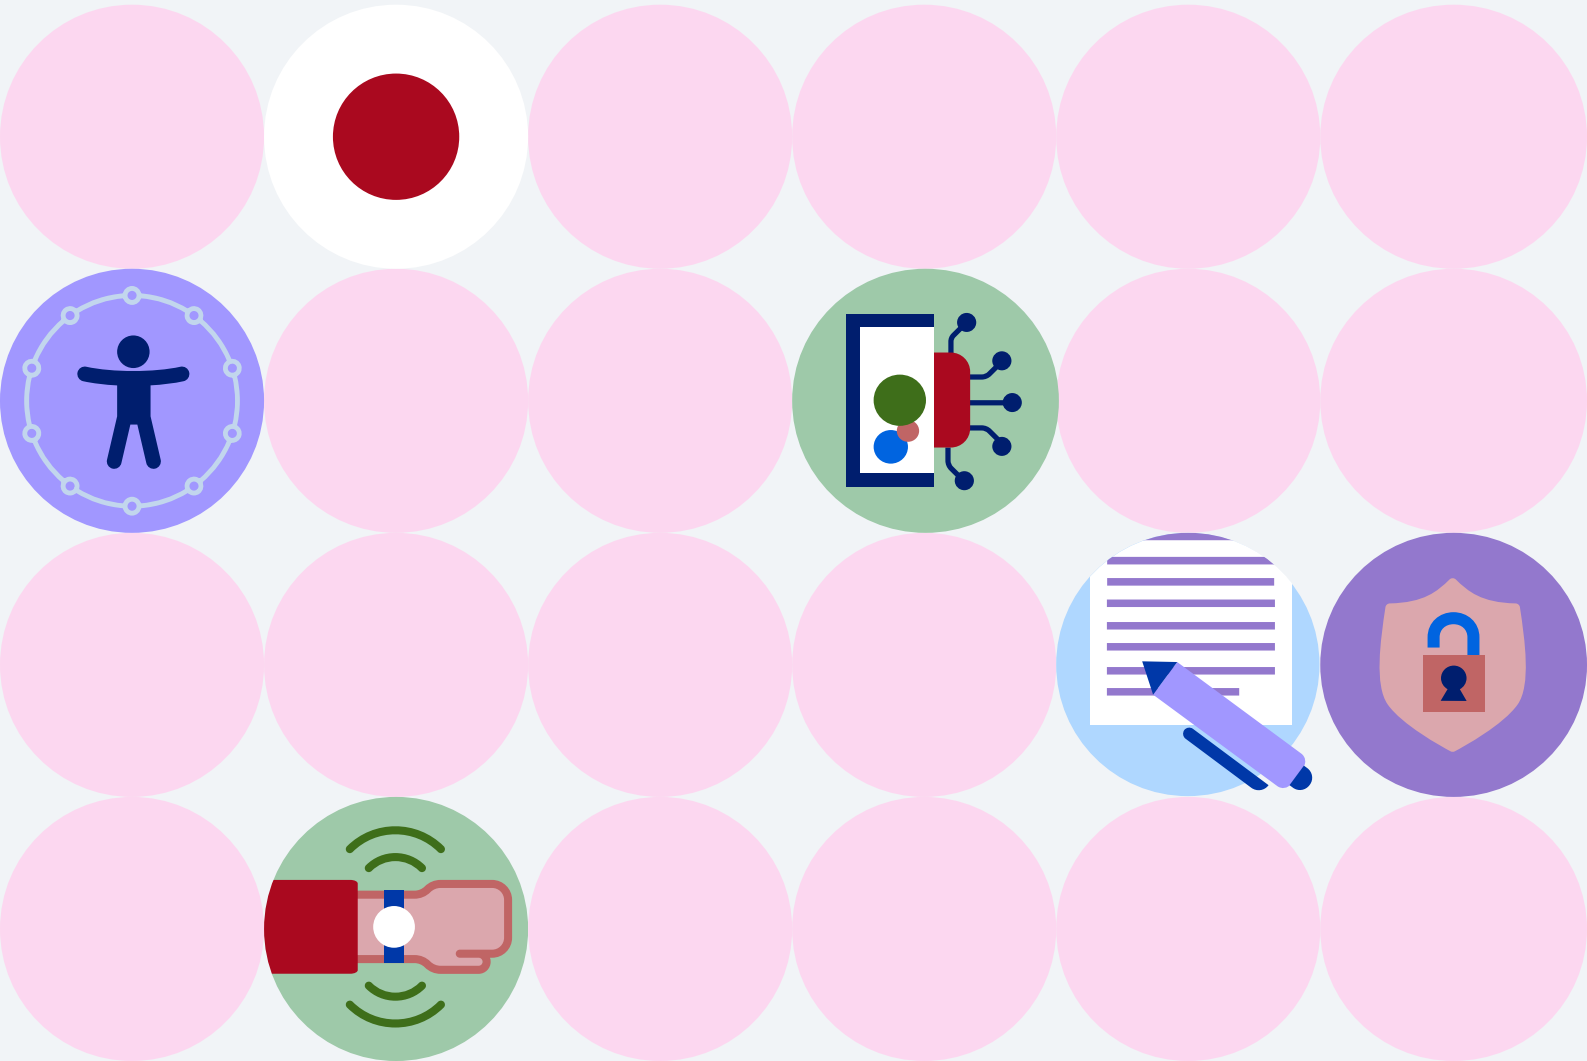
Advancing Responsible Innovation in AI and Wearable Technologies in Japan

Challenges and opportunities under the current legal guidance



Contents

About the program and this report 4

About Open Loop 4
Acknowledgements 5
Forewords 6

→ Executive summary 10

Findings & recommendations 12
Implications for Japanese policy 13

1 Chapter 1 14

1.1 Context & Objectives 15
1.2 The Policy Challenge: Governing New Data Categories in the Age of AI-Powered Wearables 15
1.3 The Goals of Open Loop Japan 16

2 Chapter 2: Findings, recommendations and opportunities 17

2.0 Introduction 18
2.0.1 About the evidence in this chapter 26
2.0.2 Our timeline and legislative developments in Japan 29
2.0.3 The central tension 29

2.1 Finding 1 – Japan's AI guidance reaches its intended audience, but participants do not find it sufficient to resolve questions of compliance with binding legal obligations 34
2.1.1 What organizations are already doing 35
2.1.2 What the data tells us 36
2.1.3 What challenges organizations are facing 37
2.1.4 What this means for the ecosystem 40
OPPORTUNITIES FOR THE PPC AND OTHER POLICYMAKERS 42

2.2 Finding 2 – Organizations have difficulty establishing how far they must de-identify personal information to use it in their AI activities 45
2.2.1 What organizations are already doing 46
2.2.2 What the data tells us 47
2.2.3 What challenges organizations are facing 47
2.2.4 What this means for the ecosystem 48
RECOMMENDATIONS FOR THE PPC 49
OPPORTUNITIES FOR OTHER POLICYMAKERS 50

2.3 Finding 3 – The classification of AI outputs and models under the APPI remains unresolved 51
2.3.1 What organizations are already doing 52
2.3.2 What the data tells us 52
2.3.3 What challenges organizations are facing 52
2.3.4 What this means for the ecosystem 55
RECOMMENDATIONS FOR THE PPC 56

2.4 Finding 4 – Cross-border rules prevent participants from using AI systems hosted outside of Japan, and cloud rules cause uncertainty and hesitancy 58
2.4.1 What organizations are already doing 59
2.4.2 What the data tells us 59
2.4.3 What challenges organizations are facing 60
2.4.4 What this means for the ecosystem 64
RECOMMENDATIONS FOR THE PPC 65
OPPORTUNITIES FOR OTHER POLICYMAKERS 67

2.5 The rules that will govern wearable devices are being written now, and the decisive questions are open 69
2.5.1 The evidence base for this finding 70
2.5.2 What participants told us 70
2.5.3 What the new category requires, and what remains open 71
2.5.4 Three questions that will determine whether the framework supports development 73
2.5.5 A non-regulatory constraint: public acceptance 75
2.5.6 What this means for the ecosystem 76
RECOMMENDATIONS FOR THE PPC 78
OPPORTUNITIES FOR OTHER POLICYMAKERS 81

3 **Chapter 3: Conclusion and next steps** 82
3.1 Summary of key takeaways 83
3.2 Implications for Japanese policy 88
3.3 Future directions 89

Annexes 91
I About Open Loop Japan 92
i.i Program Rationale 92
i.ii Partners 92
i.iii Participating Companies (The Cohort) 93
i.iv Methodology 94
i.v Limitations 97
II Key concepts & definitions 99
Endnotes 101
Bibliography 102

About Open Loop

Meta's Open Loop is a global program that connects AI experts, policymakers, and companies to help develop effective and evidence-based policies for AI and other emerging technologies by gathering detailed feedback on new or existing policies, regulations, laws, or voluntary frameworks.

The aim is to improve the quality of guidance and regulation on emerging technologies, ensuring that they are understandable, feasible in practice, and likely to achieve their intended outcome.

This work is licensed under a Creative Commons Attribution 4.0 International License.



Acknowledgements

The Open Loop Japan program was created and delivered by Meta in collaboration with The Future of Privacy Forum (FPF) and Accenture Japan.

We would like to express our gratitude to the following colleagues from FPF for their partnership and in particular their research and operational leadership on this project:

- Josh Lee Kok Thong, Managing Director for APAC
- Dominic Paulger, Deputy Director for APAC
- Bilal Mohamed, Policy Manager
- Reina Sasaki, Project Researcher
- Lauren Koek, Policy Intern
- Rachelle Faust, Senior Project Manager, Global Privacy
- Gabriela Zanfir-Fortuna, Vice President for Global Privacy

We would also like to extend our gratitude to our colleagues from Accenture for their research and operational support to this project:

- Jimmy Sia, Senior Regional Commercial Lead
- Isha Bhalla, Senior Manager
- Thomas Telandro, Senior Manager
- Kanji Endo, Consultant

Most importantly, we would like to thank the following organizations for their participation throughout the program — without their commitment and active involvement, this work would not have been possible:

Benesse Corporation

Citadel AI

Digital MICE

Hakuhodo DY Holdings

Mercari

Schematching Inc.

Techno Pro Inc.

IBM Japan, Ltd.

Shionogi & Co. Ltd.

Seven companies participated in the project on an anonymous basis.

Forewords



The convergence of advanced AI with wearable technologies is opening up entirely new possibilities for personalised experiences, preventative healthcare, and seamless digital interaction with the products and services we use every day. With these possibilities, however, come important questions about how new categories of data, including body-based and biometric data, should be collected, processed, and protected.

Getting the governance of these emerging technologies right is not a task for any single actor. It requires effective, evidence-based guidance that is clear enough to provide certainty for innovators, yet flexible enough to accommodate the rapid pace of technological development. It also requires deep collaboration between those developing the technology and those responsible for regulating it. By working together, policymakers and industry can ensure that governance frameworks are comprehensive, risk-based, implementable, and genuinely supportive of responsible innovation.

Japan is uniquely positioned to lead on these questions. Having declared its ambition to become the world's "most AI-friendly country," Japan has taken decisive steps — including the passage of the AI Promotion Act in 2025 and the recent amendments to the Act on the Protection of Personal Information (APPI) — to create a governance environment that balances innovation with effective protections. The choices made during this period will not only shape Japan's domestic AI ecosystem, but will influence regulatory approaches across the Asia-Pacific region and beyond.

It is against this backdrop that we created the Open Loop Japan program, in close partnership with the Future of Privacy Forum (FPF) and Accenture Japan, to take a structured, evidence-based approach to understanding how companies developing and deploying AI-powered wearable technologies in Japan are navigating this evolving regulatory landscape. Over the course of the program, we have worked with 16 companies operating in Japan and 20 leading AI and data protection experts to examine how the APPI and AI Guidelines for Business apply in practice to the processing of new data categories, including body-based and biometric data, and where greater clarity or coherence may be needed.

Through surveys, interviews, and expert convenings, we heard that companies developing AI-powered wearables in Japan are seeking clarity rather than latitude — practical, risk-based criteria for de-identification, confidence that trained models and their outputs are not personal data absent identifiability, cross-border transfer mechanisms that can be relied upon in advance, and workable definitional boundaries for new categories of biometric data. Where that clarity is missing, companies told us they defer or narrow viable

Forewords

use cases rather than risk non-compliance. This report offers recommendations to Japan's regulators and ministries which we believe will address these challenges. Our hope is that these insights and recommendations will be of value to policymakers and regulators as they shape the implementation of Japan's AI and data governance frameworks in the months ahead. Furthermore, we hope they are of interest to regulatory bodies across the Asia-Pacific and beyond who are grappling with similar questions about how best to govern AI and emerging data categories in ways that both protect individuals and enable innovation.

Finally, I would like to express my sincere gratitude to our partners at the Future of Privacy Forum and Accenture Japan, to the expert group who generously contributed their knowledge, and to the companies who invested significant time and energy in this effort. This work would not have been possible without your commitment to advancing responsible, innovation-enabling governance in Japan. We look forward to continuing this important work together.

Dan Hayden
Meta

Forewords



Technology development cycles are shortening and intensifying. By its very nature, technology development has usually outpaced policymaking and regulation. Yet, in fields such as AI and biometrics, the pace of development and deployment, as well as the breadth and depth of their societal implications, have brought to fore anxieties about whether current policies and regulations are fit-for-purpose. Under these conditions, the traditional distinction of roles between government, industry, academia and civil society is beginning to break down. Developing a fit-for-purpose regulatory framework therefore increasingly depends on the speed and degree with which these stakeholders can converge and collaborate to co-create policy and regulatory solutions.

It is in this spirit that the Future of Privacy Forum (FPF) took on the role of Research Partner for the Open Loop Japan program. As a global non-profit organization that advances principled and pragmatic data protection, AI and digital governance practices, FPF convenes leaders across industry, academia and the public sector to analyse, benchmark and develop best practices that support responsible innovation and regulatory compliance. We thus see Open Loop Japan as a kindred program, and a valuable opportunity to engage with relevant stakeholders, and with the broader privacy community. Equally importantly, we see it as an opportunity to advance understanding, literacy, responsible innovation and, ultimately – trust – in Japan's data protection and AI ecosystem.

The focus on Japan's policy framework is both important and timely. The present zeitgeist in Japan is defined by a distinct desire to foster a culture of responsible AI use while intentionally avoiding rigid over-regulation that could stifle technological breakthroughs. In Japan, enterprise AI adoption has jumped: from just 11% five years ago to 75% today.[1] Similarly, we see this in how Japan's AI Promotion Act – one of the first of its kind in the world – prioritises innovation with sensibility, while leaning on flexible soft law frameworks like the AI Guidelines for Business. That our work in this program proceeded amidst the passage of amendments to Japan's Act on the Protection of Personal Information (APPI) (including changes to the consent framework for AI and statistical processing) also highlights the timeliness of this project amidst ongoing efforts to position Japan as the world's most "AI-friendly country."

Yet, as any experienced policymaker would know, sound policies and regulations are a necessary, but insufficient, component for any policy objective. These policies and regulations must similarly be reflected in the lived realities on the ground. This is where we hope FPF's objective and data-driven research work – comprising surveys, interviews with

[1] https://www.bbc.com/news/articles/cwymw4434v7o

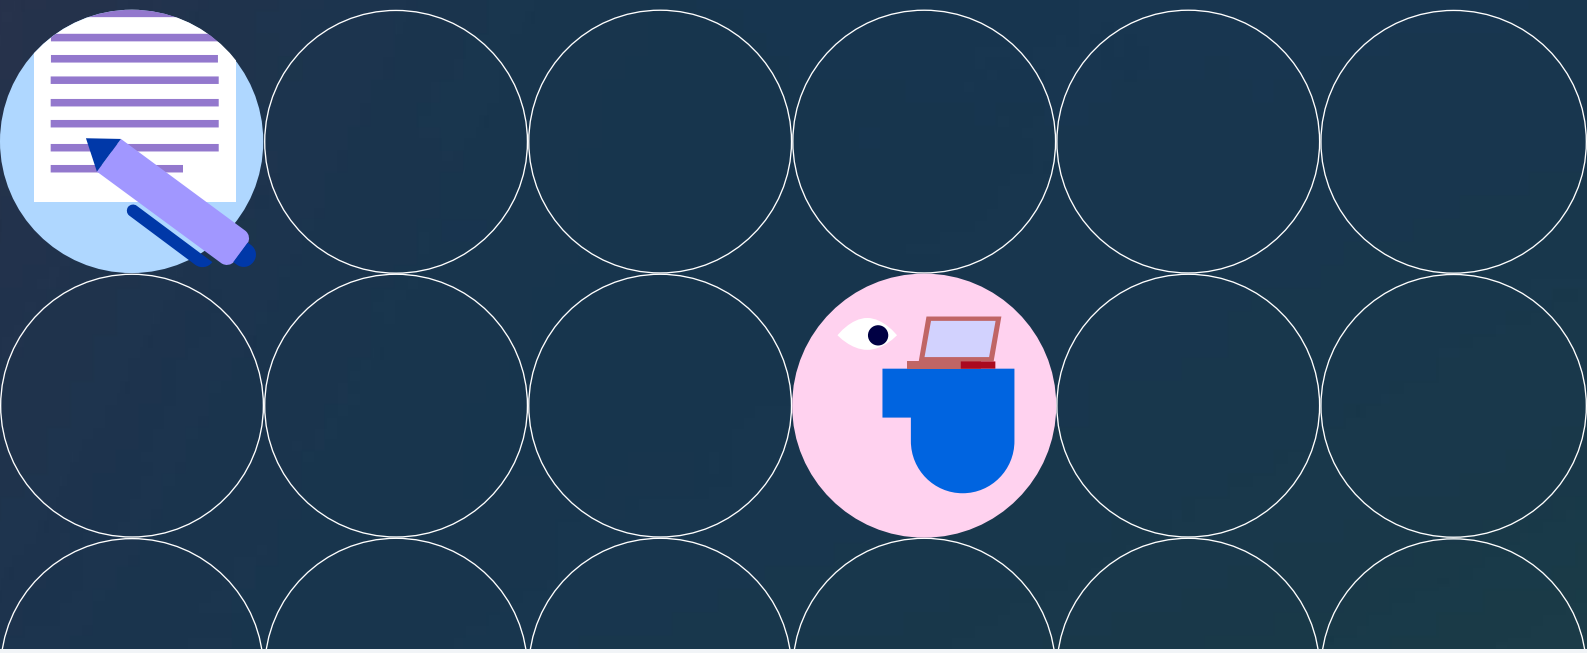
Forewords

companies, and analysis of primary data – can contribute to a fuller understanding of the application and intersection of Japan's AI and data frameworks on the ground. We approach this work with the deepest respect for all work that has come prior, and are humbled to be able to contribute our part towards the vision of the innovative and trustworthy use of AI in Japan.

It leaves me to express our sincere gratitude to all the participating companies, members of the Expert Group, our consortium partners Meta and Accenture, and all who have contributed to this work in one way or other. We hope this report serves as a practical reference for policymakers and practitioners as Japan's AI and data landscape continues to take shape, and look forward to continuing being a trustworthy partner and resource to all in Japan.

Josh Lee Kok Thong
Managing Director, Asia-Pacific, Future of Privacy Forum

Executive summary



Open Loop Japan set out to build an evidence base on how Japan's AI and data governance frameworks operate in practice, and where greater clarity would allow them to support both data protection and responsible innovation, at a time when Japan aims to become the world's "most AI-friendly country." The program focused on how organizations developing and deploying advanced AI systems in Japan (including AI-powered wearable devices) navigate requirements under Japan's three primary AI governance frameworks:

- **The Act on the Protection of Personal Information (APPI)** – a binding law administered by the Personal Information Protection Commission (PPC);
- **The Act on the Promotion of Research and Development, and Utilization of AI-related Technologies (AI Promotion Act)**, which sets a strategic framework for promoting AI use in Japan but contains limited obligations for private-sector actors; and
- The **AI Guidelines for Business (Version 1.2)**, drafted by the Ministry of Economy, Trade and Industry (**METI**) and Ministry of Internal Affairs and Communications (**MIC**), set out voluntary guidance.

Sixteen companies operating in Japan and representing a broad range of sectors, completed a structured bilingual survey between June and August 2026; eight took part in follow-up interviews. This evidence base was supplemented with written input from four members of the program's Expert Group with relevant specialized knowledge – one industry professional, one academic, and two legal practitioners based in Japan – on questions relating to wearable devices.

The following observations run across the research:

1. **Regulatory uncertainty is not preventing AI adoption, but it is shaping it.** It produces internal review costs, conservative defaults, and use cases that are narrowed or never attempted. It also leaves organizations to reach different conclusions about the same requirement, with no visibility for the regulator into where they differ.

2. **Participants are not asking for lighter obligations but rather, for rules they can apply before they act:** criteria assessable in advance, examples, model clauses, and checklists.

3. **75% of participants anticipate that the 2026 amendments to the APPI will constructively reduce barriers to using personal information for AI** – the largest measure of agreement in our survey.

Findings & recommendations

1 Japan's AI guidance is widely consulted, but does not resolve questions of APPI compliance

11 of 16 participants (69%) take reference from the AI Guidelines for Business, yet only one (6%) finds it sufficient to establish what the APPI requires without further legal advice. METI and MIC could supply examples, model clauses, checklists, and an explicit map from the Guidelines to APPI obligations. Together with the PPC, METI and MIC could establish a sandbox, structured dialogue pathways, and published case studies of compliant deployment.

2 Organizations cannot establish in advance how far they must de-identify data

Practical, risk-based anonymization standards are the cohort's highest-ranked policy ask (10 of 16, 62%). The PPC should publish risk-calibrated criteria that organizations can assess for themselves, and state the legal position secured once they are met, including under the anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報) regimes.

3 The classification of AI models and outputs under the APPI remains unresolved

Half of participants cite the risk that outputs reconstruct or infer personal information. The PPC should consider confirming that a trained model is not, of itself, a personal information database (個人情報データベース等) under the APPI – consistent with its existing interpretation for search engines and with the premise of the new statistical creation exemption – and issue guidance on proportionate measures against output inference risk.

4 Cross-border transfer mechanisms are not determinable in advance

7 of 16 (44%) are uncertain which overseas providers or jurisdictions satisfy the APPI, and 11 of 16 (69%) still seek confirmation on the effect of Global Cross-Border Privacy Rules (CBPR) certification, notwithstanding a clarification on the effect of this given by the PPC in December 2025. This is a question of visibility rather than of law. The PPC should publicize its clarification on the effect of Global CBPR certification, give certification a defined benefit, state the legal effect of its forthcoming model clauses, and draft those clauses, and restate the cloud exception, with AI processing patterns in view. METI and MIC could align the Contract Guidelines on Utilization of AI and Data with those clauses.

5 The rules that will govern wearable devices are being written now

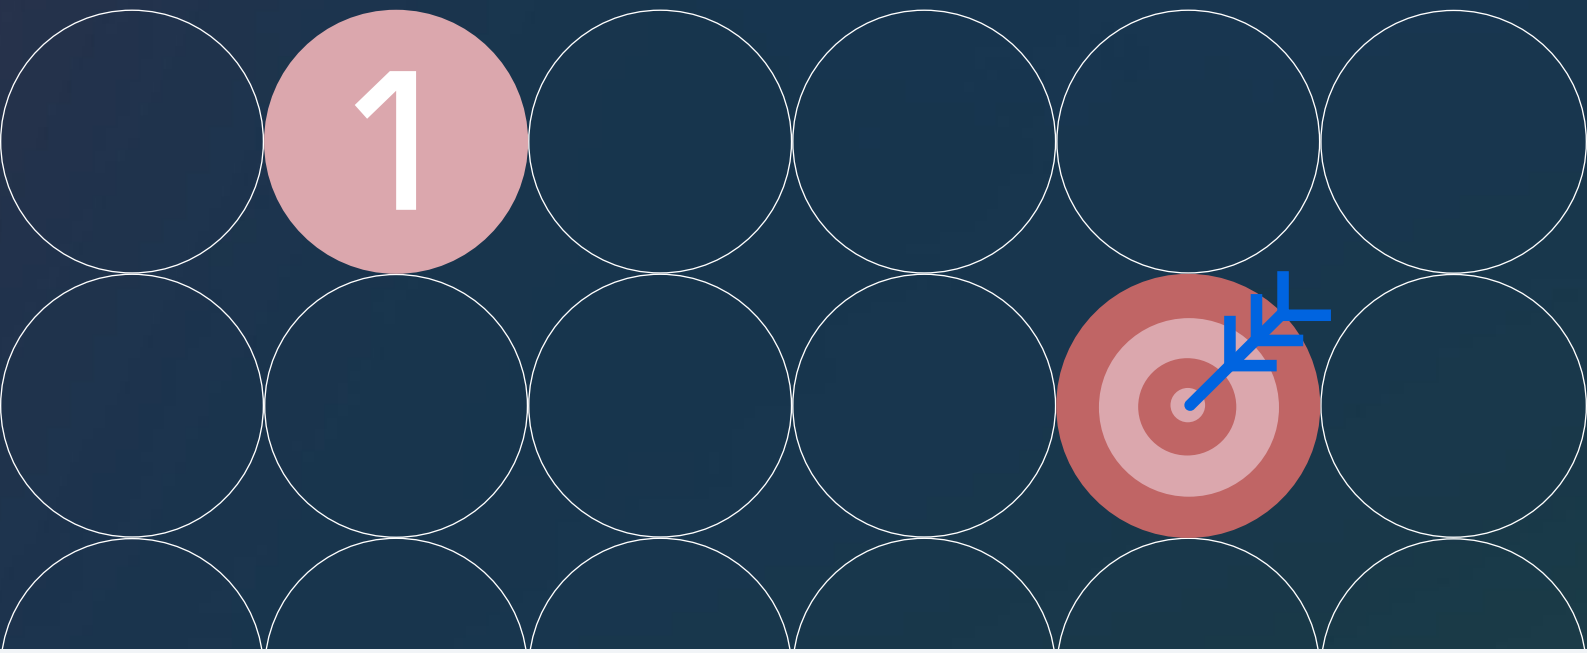
During the commencement period, the PPC should confirm that a person who chooses to wear a device, and is told what it collects, is easily aware of that collection; clarify whether processing to identify an individual is decisive; state which modalities are in scope; and address whether notice of recording extends to later recognition. Throughout, guidance should specify conditions rather than mechanisms, and the PPC and METI could set the conditions under which incidental capture of bystanders is adequately notified.

Implications for Japanese policy

Almost nothing participants ask for requires further legislation. The questions identified are questions of interpretation, which Cabinet Orders and guidance can answer within the period before the amended APPI takes effect. That interval is the operative period: guidance published early enough can be built to, whereas guidance arriving shortly before commencement cannot. Several questions also cut across institutional boundaries, and would benefit from the PPC, METI and MIC acting together. For the new APPI category of "specific biometric personal information" in particular, the principal risk is not under-regulation, but premature over-breadth that forecloses the use cases the category was created to govern.

CHAPTER ONE

Chapter 1



1.1 Context & objectives

Japan's Evolving AI and Data Governance Framework

Japan has positioned itself as one of the world's most ambitious adopters of artificial intelligence, with a stated national goal of becoming the "most AI-friendly country." This ambition is reflected in a series of significant regulatory developments: the passage of the AI Promotion Act in May 2025, which establishes a strategic framework for AI development and deployment; the ongoing triennial review of the Act on the Protection of Personal Information (APPI), which is examining how Japan's data protection framework should evolve to facilitate new technological developments; and the continued refinement of the AI Guidelines for Business, a living document which provides guidance to help companies to achieve responsible AI use. Japan's leadership on AI governance extends beyond its borders — as chair of the G7 Hiroshima AI Process, Japan has played an influential role in shaping international approaches to AI governance, and the regulatory choices made domestically will carry significant weight across the Asia-Pacific region.

1.2 The Policy Challenge: Governing New Data Categories in the Age of AI-Powered Wearables

At the intersection of these regulatory developments lies a pressing and under-examined policy challenge: how should existing and emerging governance frameworks apply to the new categories of data generated by AI-powered wearable technologies? Wearable devices from smartwatches and fitness trackers to AI-enabled smart glasses may require the collection and processing of body-based and biometric data in order to deliver the services requested by their users. As AI capabilities enable more sophisticated and useful inferences to be drawn from this data, questions arise about the adequacy of current definitions, processing grounds, and safeguards which may be effectively applied. Companies developing and deploying these technologies in Japan face growing uncertainty about how relevant obligations, including the treatment of Personally Referable Information (PRI), external transmission rules, and sensitive data categories apply to their products and services. Without clear, pragmatic, and innovation-enabling guidance, there is a risk that this regulatory uncertainty could inadvertently stifle innovation in a sector where Japan has both significant commercial potential and the opportunity to set global standards.

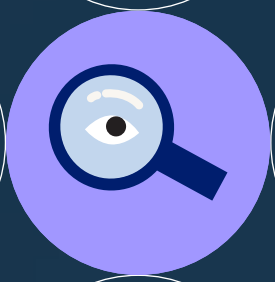
1.3 The Goals of Open Loop Japan

Open Loop Japan was designed to address this challenge through structured, multi-stakeholder collaboration. By bringing together companies working across the AI and wearable technology ecosystem with policymakers, privacy experts, and academics, the program generated evidence-based insights that can inform the practical development and implementation of Japan's governance frameworks. The program has three core objectives: first, to understand how companies are currently interpreting and implementing regulatory requirements as they relate to new data categories, including body-based and biometric data; second, to identify where existing frameworks provide sufficient clarity and where gaps, ambiguities, or friction points remain; and third, to develop actionable, evidence-based policy recommendations that support the dual goals of data protection and continued innovation. By bridging the gap between policy and practice, Open Loop Japan seeks to ensure that Japan's regulatory approach enables — rather than inadvertently stifles — the responsible development and deployment of AI-powered wearable technologies.

CHAPTER TWO

Findings, Recommendations and Opportunities

2



2.0 Introduction

In Japan, AI governance relies on the intersection of the Act on the Protection of Personal Information (**APPI**), the Act on Promotion of Research and Development, and Utilization of AI-related Technology (**AI Promotion Act**), and AI-specific voluntary guidelines, such as the AI Guidelines for Business (Version 1.2).

Although the APPI is technology-neutral, it imposes binding obligations regarding the use of personal information; this is highly relevant to AI because the unstructured datasets used to train and refine generative AI models frequently contain personal information, and because personal information may be disclosed directly to AI systems. Furthermore, the 2026 amendments to the APPI have reinforced the Personal Information Protection Commission's (PPC) role in AI governance: the Cabinet's summary of the Amendment Bill states that a new exemption from the consent requirements, directed at the creation of statistical information, extends to AI development that can be characterized as statistical creation.

The AI Promotion Act establishes a legal framework for AI systems in Japan, but it provides minimal obligations for private-sector actors and carries no penalties.

Consequently, actionable obligations for private sector actors are primarily found in the APPI – administered by the PPC – and voluntary frameworks, such as the AI Guidelines for Business (Version 1.2), which are jointly drafted by the Ministry of Economy, Trade and Industry (**METI**) and the Ministry of Internal Affairs and Communications (**MIC**).

These frameworks vary significantly in their regulatory scope, enforcement mechanisms, and underlying policy objectives. Nevertheless, any organization seeking to develop or deploy AI in Japan must consider all three frameworks simultaneously. **It is therefore critical that they closely align to avoid regulatory friction that could hinder the ability of organizations to innovate and compete in the AI sector.**

This chapter sets out what participants told us about working at the intersection of these three frameworks. It identifies where those frameworks are working as intended, where participants reached different conclusions about what is expected of them, and how the PPC, METI, MIC, and other policymaking bodies in Japan could support a more consistent reading of the requirements and guidance.

Below, we set out the five key findings based on participants' responses to a survey, and follow-up interviews. A detailed account of the research methodology is provided in **§2.0.1**. This is followed by a breakdown of the results that lead to each of the five key findings (**§2.1-2.5**). Each finding closes with the related Recommendations and/or Opportunities for the policymakers in Japan.

FINDING 1 (§2.1)

Japan's AI guidance reaches its intended audience, but participants do not find it sufficient to resolve questions of compliance with binding legal obligations

Participating companies report that they do consult voluntary AI guidance but encounter challenges implementing this guidance and do not find the guidance sufficient to determine questions of compliance with binding legal obligations relating to AI, such as those in the APPI, without the need for further legal interpretation. 11 of 16 participants (69%) reported consulting the AI Guidelines for Business (Version 1.2), but only one (6%) found it clear enough to determine what the APPI requires without further legal advice.

Participants accept the principles on which Japan's AI guidance is based but face difficulties translating them into practice, or identifying whether a particular vendor arrangement is acceptable. Participants were of the view that they would benefit from model clauses, checklists, and examples they can apply to the cases in front of them.

FINDING 1 CONTINUED

OPPORTUNITY 1

For METI and MIC:

> **Supply practical resources in the next iteration of the AI Guidelines for Business**

These include examples, model contract clauses, standard review checklists and sector annexes, together with an explicit map from the Guidelines to the APPI obligations they engage.

OPPORTUNITY 2

For METI, MIC, the PPC and cross-government:

> **Establish a sandbox and publish case studies of compliant AI deployment in Japan**

Organizations have specific questions but no established channels to seek answers. A sandbox allows approaches to be tested before rollout; dialogue pathways allow a particular question to be put to a regulator; case studies give a reference point no organization can produce alone.

FINDING 2 (§2.2)

Organizations have difficulty establishing how far they must de-identify personal information to use it in their AI activities

A third of the cohort uses anonymized or aggregated data in their AI activities. However, participants cited difficulty establishing in advance whether their de-identification efforts satisfy the APPI's threshold for anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報) before their data processing falls outside the APPI's scope.

RECOMMENDATION 1

Publish practical, risk-based criteria for how far an organization must de-identify personal information before they can use it in AI-related activities

Criteria should be calibrated to the risk of the use case, assessable by the organization itself before it acts, and should state what legal position is secured once they are met. Guidelines should include examples that show when data has been successfully anonymized, and when it has not.

OPPORTUNITY 3

For METI, MIC, the PPC, and standards bodies:

Establish external assurance for de-identification tooling

Participants cannot verify that the tools they use to de-identify data work as intended, and no third-party body exists to assess them against published criteria. A safe harbour for organizations applying published criteria in good faith (i.e. through documented and thorough risk assessments) would provide that assurance.

FINDING 3 (§2.3)

The classification of AI outputs and models under the APPI remains unresolved

Half of participants (50%) reported that Japan's existing AI governance frameworks have yet to settle the obligations that apply to personal information entered into an AI model, information the model returns, or the trained model itself. Organizations reach different conclusions, and one participant reports proceeding without having settled the question.

RECOMMENDATION 2

> **Consider issuing guidance clarifying that trained models and their outputs are generally not personal data absent identifiability and consider issuing guidance on minimizing output inference risk**

Confirm whether a trained model constitutes a "personal information database" under the APPI, and where in a development chain any resulting obligation sits. Separately, set out the measures organizations can take against the risk that outputs reveal personal information – input controls, testing before deployment, and retention practices for data used in fine-tuning.

FINDING 4 (§2.4)

Cross-border rules prevent participants from using AI systems hosted outside of Japan, and cloud rules cause uncertainty and hesitancy

Just under half of the cohort reports uncertainty about which overseas providers and jurisdictions satisfy the APPI's requirements for equivalent protection and equivalent measures. Participants cite practical challenges mapping data flows and negotiating terms for transfer agreements with AI service providers.

RECOMMENDATION 3

> **Make the cross-border transfer mechanisms determinable in advance**

Publicize the PPC's existing clarifications on the effect of Global CBPR certification, and give this certification a defined benefit.

State the legal effect of adopting the PPC's forthcoming model clauses for cross-border data transfers, draft them with AI processing patterns in view, and restate the cloud exception against those patterns.

FINDING 4 CONTINUED

OPPORTUNITY 4

For METI and MIC:

> **Align the Contract Guidelines on Utilization of AI and Data with the PPC's (forthcoming) model clauses**

The PPC's forthcoming model clauses will address what an exporter must secure from an overseas recipient of personal information. METI's Contract Guidelines already address commercial contracting for AI and data. Aligning the two as the clauses are drafted would let a single contract satisfy both.

OPPORTUNITY 5

For the PPC and METI, with industry:

> **Make provider adoption of the (forthcoming) model clauses visible, so exporters can verify rather than negotiate**

Publishing the clauses in English alongside Japanese, and maintaining a public record of providers that have adopted them unamended, would turn an exporter's task into checking a list rather than opening a negotiation it may lack the leverage to win.

FINDING 5 (§2.5)

The rules that will govern wearable devices are being written now, and the decisive questions are open

Only one participant develops or deploys advanced wearable devices for the Japanese market, and none is a device manufacturer. That absence sits alongside sustained demand for guidance on the new category of "specific biometric personal information" (特定生体個人情報)[2] - added in the 2026 amendments to the APPI. This new category is scoped by the conditions by which data is captured, rather than by the purpose for which the data is processed. How that scoping applies – in particular to a device's own user, as distinct from people captured incidentally – will be determined by Cabinet order.

Organizations appear to be seeking clarity on this in advance of building.

RECOMMENDATION 4

> **Provide risk-based and workable definitional guidance on "specific biometric personal information" during the commencement period**

Confirm that a person who chooses to wear a device, and is told what it collects, is easily aware of that collection. Clarify whether processing in order to identify an individual is a decisive factor. State which modalities are within scope and how boundary cases will be assessed as technology changes. Address whether a notice given for recording extends to later application of recognition technology. Throughout, specify the conditions to be satisfied rather than the means of satisfying them.

[2] Article 16(5) of the amended APPI, "specific biometric personal information" (特定生体個人情報) is defined as personal information that contains a "specific biometric individual identification code" (特定生体個人識別符号). The latter is defined as data created by converting a physical characteristic that can be obtained without special technology or high costs, and without the individual easily recognizing that it is being collected. This report uses the term "specific biometric personal information" (特定生体個人情報) throughout because the amended APPI attaches primary compliance obligations directly to this broader category. These obligations include notifying individuals prior to handling of such information and complying with requests to suspend the use of such information, or provision of such information to third parties.

FINDING 5 CONTINUED

OPPORTUNITY 6

For the PPC and METI, in the commencement period:

> **Establish the conditions under which incidental capture is adequately notified**

Where a device records people who are not its user, the organization cannot easily contact them. Rather than prescribing mechanisms, state the conditions a solution must meet: that a person nearby can perceive that collection is occurring, can obtain further information, and has a means of objecting.

2.0.1 About the evidence in this chapter

This chapter is based on:



A structured bilingual survey completed by **16 participants** between June and August 2026



Follow-up interviews conducted in Japanese with **8 of the participants** throughout August 2026

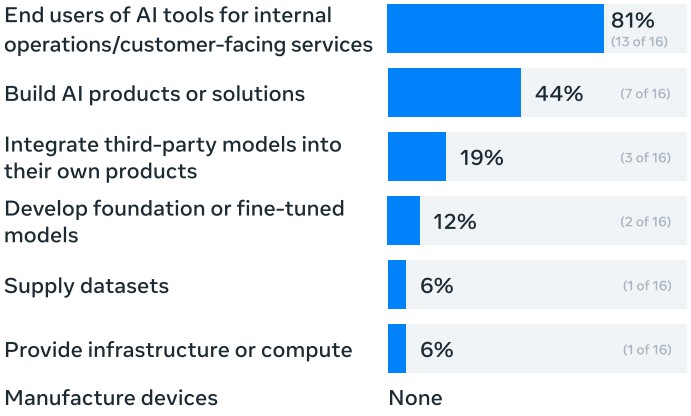
Taken together, the cohort of 16 participants spans a diverse range of sectors, including consumer health, education, e-commerce, energy, food and life-science manufacturing, electrical equipment manufacturing, information technology, pharmaceuticals, human-

resources services, advertising, semiconductors, and consulting. We present all findings in aggregate and identify no individual respondent or organization in this report.

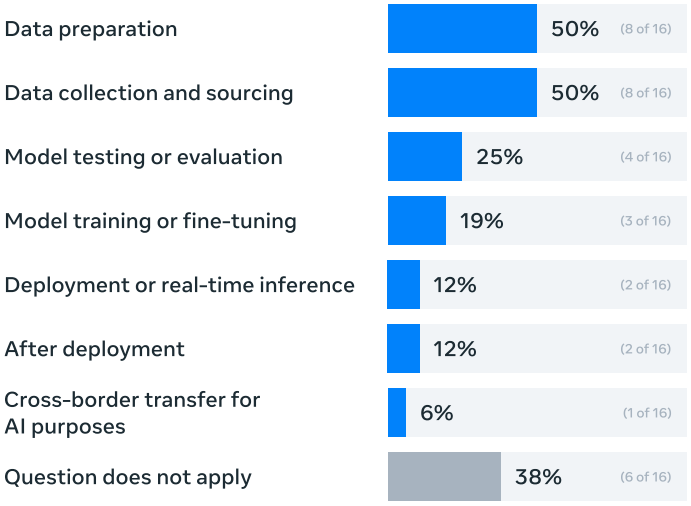
Two features of the cohort shape the findings that follow:

First, this cohort adopts AI more than it builds AI

Participants described their roles as follows (note that participants were permitted to choose multiple options):



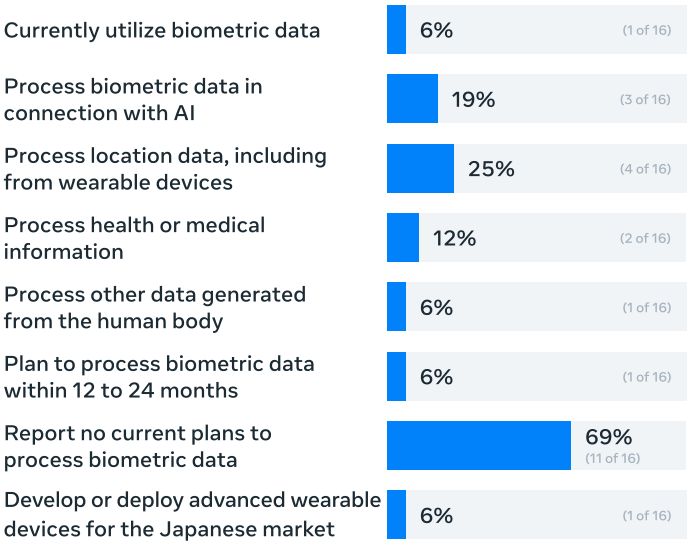
Participants also tend to handle personal information at the earliest stages of their AI workflows:



5 of 16 participants (31%) indicated that they process no personal information in connection with AI at all.

The challenges that this cohort faces are therefore concentrated at the point where organizations collect and prepare data for use in third-party AI tools, and in the governance decisions that precede the deployment of these tools, rather than in the mechanics of training AI models.

Second, the cohort's engagement with biometric data[3] and data collected by wearable devices,[4] is limited, though not absent





A note on terminology regarding participants' preferences for policy reform

The survey asked two separate questions about policy preferences:

1. One question asked participants to rank the policy changes and support mechanisms that would most benefit their AI operations. This chapter refers to the responses to this question as **"ranked policy asks"**.
2. A separate question asked which priorities Japan's regulatory framework should adopt more generally, with a multiple selection response format offering differently worded options. This chapter refers to responses to this question as "**framework priorities**".

The distinction matters because the two questions focus on different preferences: "ranked policy asks" reflect what a participant **wants for its own operations**, in order of importance to that organization, while "framework priorities" record **what a participant thinks that Japan's framework needs**, whether or not it affects them. A participant may accordingly rank an issue more highly under one question than the other.

[3] In the survey, we defined **"biometric data"** as data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of an individual, which allow for that individual to be uniquely identified. Under the APPI, this falls under the category of an "Individual Identification Code" (個人識別符号), which includes data converted from physical characteristics for computer-based recognition (e.g., facial recognition data, fingerprints) (APPI, Article 2 and relevant Cabinet Orders).

[4] In the survey, we defined **"wearable devices"** as electronic devices designed to be worn on the user's body or incorporated into clothing (e.g., smart glasses, AR/VR headsets, smartwatches, and health trackers). These devices are frequently equipped with advanced sensors (such as cameras, microphones, and accelerometers) and edge computing capabilities, enabling the continuous collection and processing of both user-specific body-based data and ambient environmental data.

2.0.2 Our timeline and legislative developments in Japan

We began collecting data while a bill to amend the APPI was still before the National Diet: the Bill for the Act Partially Amending the Act on the Protection of Personal Information and Other Related Acts (Cabinet Bill No. 54 of the 221st Diet) (the "**Amendment Bill**"). That Amendment Bill, which implements proposals from the PPC's triennial review of the APPI, was approved by the Cabinet and submitted to the Diet in April 2026. Our survey, which commenced in mid-June 2026, therefore asked participants about the proposed amendments. During the course of the survey, the Diet passed the Amendment Bill, and it was promulgated on 17 July 2026. Most of the new obligations – save for specific "early enforcement" provisions on penalties – are scheduled to take effect by 16 July 2028.

Where participants expressed expectations about the reform during the survey, those were expectations about the Amendment Bill, rather than assessments of enacted law, and we report them as such. This does not materially affect the findings of the study, as the Amendment Bill as enacted is substantially the same as the version tabled in April 2026. Nevertheless, 7 of 16 participants (44%) still reported taking reference from the amendment proposals, even while the Amendment Bill was still before the Diet.

Interviews conducted after the Amendment Bill had passed gave us an opportunity to revisit this. Participants' views appeared largely unchanged: attention had shifted from the primary legislation to the subordinate rules and guidance that will determine how it applies, and one participant that supplies AI services to other organizations in Japan reported no discernible change yet in what clients raise with it.

2.0.3 The central tension

Participants expect the reform of the APPI to help them. 12 of 16 (75%) anticipated that the amendments would constructively reduce regulatory barriers to using personal information for the development and deployment of AI – the largest measure of agreement recorded anywhere in our survey. This chapter therefore arrives at a moment when industry is receptive to reform already under way, and our recommendations are framed accordingly: they propose ways to give effect to that reform in practice, rather than arguing for something different.

That expectation sits alongside a markedly less positive assessment of the framework as it currently stands. Participants do not find the current framework clear.

- Asked whether the APPI's requirements for training AI models are clear and unambiguous, 8 of 16 (50%) disagreed, and only 3 of 16 (19%) agreed.
- Asked whether the AI Promotion Act and voluntary guidelines together set out regulatory expectations clearly, 9 of 16 (56%) disagreed, and again, only 3 of 16 (19%) agreed.
- On the usefulness of the voluntary guidelines for navigating APPI compliance, only one participant of 16 (6%) found the guidelines sufficiently clear for compliance purposes.

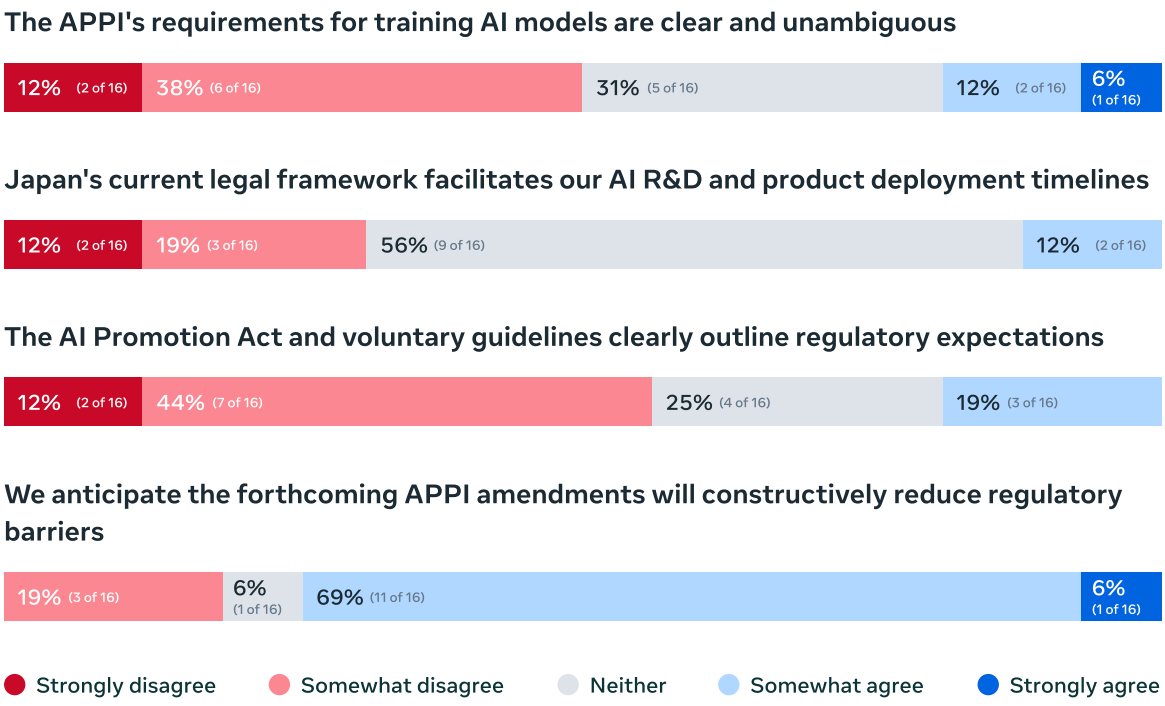
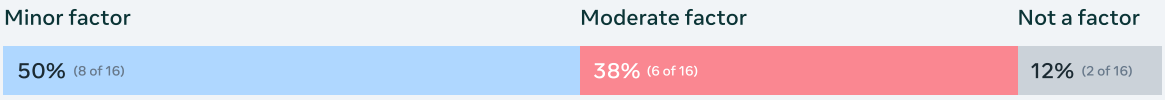


FIGURE 2.1
Perceptions of the current legal landscape

Nevertheless, this regulatory uncertainty is not preventing organizations from developing and deploying AI. Participants were asked to rate, on a five-point scale, how far regulatory uncertainty around the APPI or the AI Promotion Act affected their AI development or deployment timelines.

Notably, no participant considered regulatory uncertainty to be a "significant" or "primary factor." Rather, participants' responses were as follows:



This evidence suggests that for this cohort, regulatory uncertainty does not lead to product cancellations or abandoned deployments. Rather, it manifests as **increased internal review costs, conservative default settings, decisions taken without a settled view as to their legality, and friction moving from proofs of concept to actual deployment.**

Two participants describe the compliance review burden directly

- One observes that even a single AI governance question required consideration of the APPI, AI-related guidelines, cloud and vendor management, information security and contracting practices, which places a heavy load on reviewers and applicants alike. What the business needs, in that participant's view, is to be able to judge quickly and consistently whether a given AI tool is safe, whether staff may enter personal or confidential information into it, and how the rules treat transfer to an overseas cloud service or an AI model hosted overseas.
- A second reports ambiguity over whether their data falls within the APPI's definition of personal information, whether a proposed use goes beyond the purpose specified when it was collected, how cross-border transfer requirements apply, and where responsibility lies between the parties when an AI system makes a decision.
- A third describes the scale of this work, and where the difficulty lies within it. It conducts an estimated 300 to 400 assessments of cloud services a year, of which between a third and three-quarters involve AI. The participant specifically highlighted the overlap between the APPI, AI-related guidelines, cloud and vendor management, and information security and contracting practice as a challenge. All of these frameworks bear on the same decision, and reviewers consequently have to check the same ground more than once. The participant explains that it has produced an internal flowchart but that this cannot be applied universally – each answer depends on the use case, the data involved, the services in use and the combination of those factors, and it has not been able to systematize them. It added that its reviewers sometimes question whether their own skills are sufficient for the task, and that because stopping the reviews would bring the business to a standstill, the organization continues as best it can.

Two participants specifically referred to internal approval timelines

- Asked how long internal approval takes for an AI use case involving personal information, one participant estimated two weeks to a month, occasionally longer, across three functions: (1) security, (2) privacy, and (3) AI governance.

- A second described its own approval process as long, measured in months, with some proposals abandoned outright, and others where approval proves unattainable however long is spent. The same participant noted that its review capacity is already committed to compliance checks under other legislation (it gave export control as an example), and that adding further personal data checks on top would be a substantial burden.

One participant highlighted the **importance of in-house capabilities**

(Describing a case in which having in-house expertise avoided what would otherwise have been six months of wasted work)

- The participant described being approached by a data broker proposing a common identifier solution that the data broker presented as compliant with the APPI. The participant assessed the solution and concluded that it could not be used in Japan, recommending that the collaboration be dropped. Approximately six months later, it became clear to the market that the solution could not be used. The participant described the guidance as merely listing examples that are clearly acceptable and clearly unacceptable, but offering little on the cases in between, not indicating what to be careful about when assessing them. Rather, the judgment rested instead on the participant's own familiarity with the area, the views of other professionals in the field, and regulatory approaches in other jurisdictions. The participant considered that if a colleague without that background had been asked to consider the same question, then the participant’s organization would likely have wasted six months on collaboration negotiations, contract drafting, and development work that would ultimately have been discarded.

Two participants specifically raised concerns about adopting AI before the legal position is settled

- One participant reports that uncertainty under the APPI is leading to business designs built on optimistic assumptions about how the regulatory framework will be interpreted. The participant considers this a business risk. The interviews elaborated on this, citing past industry controversies where business models – such as employment platform tracking and data broker initiatives – were designed on the assumption that cookie-linked data did not constitute personal information, only to face severe regulatory and public backlash post-deployment.

- A second participant, describing internal use of a fine-tuned generative model, reports that it cannot tell how the APPI treats personal information entered into the model, or whether what the model returns counts as personal information. The participant reported that it has continued to use the system without settling the question, and considers that the wider debate has given far too little attention to whether proceeding on that basis is acceptable.
  The same participant made this point directly in its survey response, arguing that relief under the APPI needs to be considered separately for three things: developing AI models, using AI systems, and handling personal information arising from having used AI. On its account, the reform addresses the first, while the second and third have attracted less attention.
- A third participant described the position from the perspective of a smaller organization. In its sector, which involves substantial front-line operations, it identified multiple potential efficiency gains from AI, including multilingual customer support and operational management. However, uncertainty as to the legal position obliges it to be cautious in deciding whether to adopt AI. It also observed that for small and medium-sized organizations, the cost of working out compliance operates as a barrier to entry in practice.

Nevertheless, proceeding with AI adoption on unresolved legal bases leaves the organization exposed, leaves individuals with no reliable route to exercise their rights, and denies the regulator visibility into how industry practice has moved. The way forward, then, is greater clarity on what the existing rules require – clarity that, where the amendments were intended to enable AI innovation, resolves ambiguities in a workable, innovation-enabling direction (for example, confirming that properly de-identified data and model outputs fall outside the APPI's scope), rather than loosening or tightening of the rules themselves.

2.1 Finding 1

Japan's AI guidance reaches its intended audience, but participants do not find it sufficient to resolve questions of compliance with binding legal obligations

As the AI Promotion Act deliberately avoids heavy, prescriptive mandates and instead prioritizes an "innovation-first" blueprint, Japan's voluntary AI guidance – such as the AI Guidelines for Business (Version 1.2) – continues to play an important role in shaping organizations' governance of their AI activities and interpretation of other binding frameworks, such as the APPI. In this context, Japan's guidelines have achieved something rare and phenomenal: organizations genuinely use them.



(11 of 16 participants)

69%

take reference from the AI Guidelines for Business (Version 1.2)



(3 of 16 participants)

19%

treat these guidelines as the single most impactful framework for their day-to-day AI operations



(1 of 16 participants)

6%

find these sufficient, without the need for further legal interpretation, to establish what laws like the APPI require of its AI activities

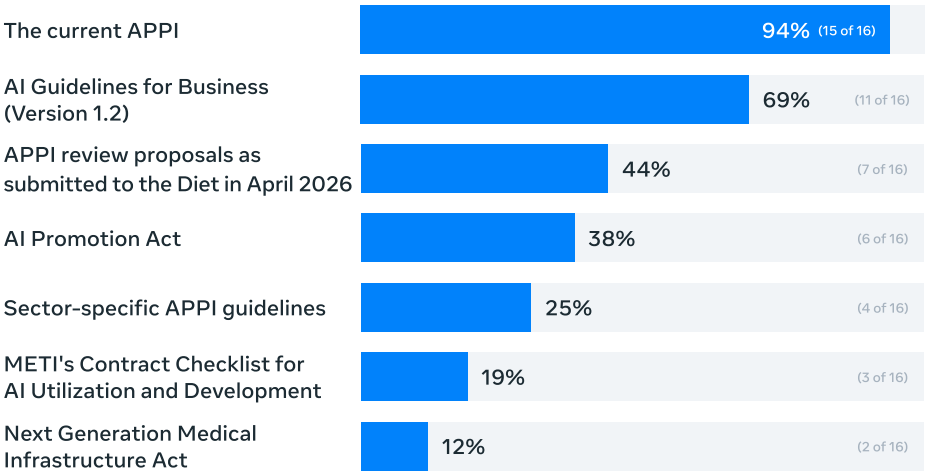
Where guidance states principles rather than rules, the work of converting those principles into something a reviewer can apply falls to each organization separately – and participants pay for it in internal guideline-writing, checklist-building and review time.

2.1.1 What organizations are already doing

Participants were asked which domestic frameworks they take reference from in their AI activities. The list spans instruments of very different legal status:

1. Binding laws, such as the APPI
2. Interpretative guidance issued by regulators such as the PPC
3. Voluntary guidelines that carry no legal force but could be relied on for interpretation purposes

Participants were permitted to choose multiple options. Their responses are as follows:

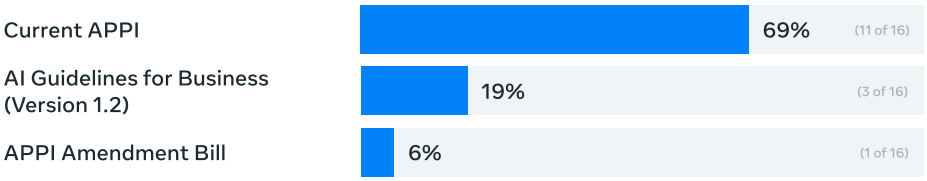


Reference to the APPI is near-universal, at 15 of 16 (94%), which reflects its status as the binding framework governing personal information in AI activities. The AI Guidelines for Business (Version 1.2) follow at 11 of 16 (69%), despite carrying no standalone binding legal force. Notably, the AI Promotion Act, though technically binding, is referenced by 6 of 16 (38%) – fewer than the voluntary guidance issued under the broader policy framework it establishes.

That may be unsurprising. The AI Promotion Act sets out the government's approach to promoting the development and use of AI and places only non-binding "best-effort" duties on private-sector actors, with no direct penalties. Its enforcement is indirect: the government must publish guidelines (Article 13) and may investigate cases where AI infringes citizens' rights and issue guidance to businesses.

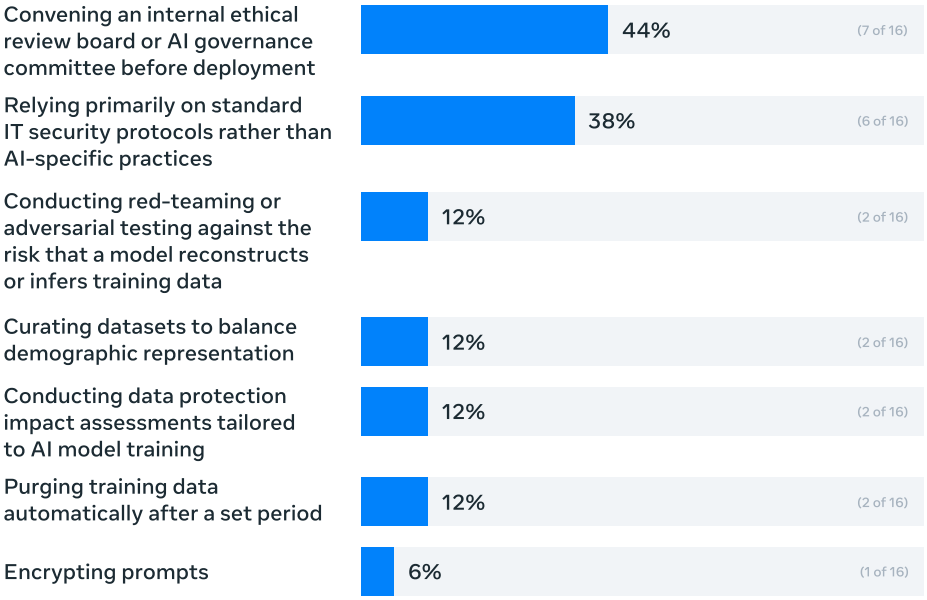
For day-to-day AI activities, however, the AI Promotion Act offers organizations relatively few concrete rules to take reference from. The practical position is therefore a division of labor between two instruments of very different status: the legally binding obligations that bear on AI activities sit in the APPI, while expectations about how those activities should be conducted are articulated in the voluntary guidance.

Participants were also asked which single framework was most impactful to their day-to-day AI operations:



For some participants – particularly those providing products or services internationally – foreign frameworks also serve as the operative standard. One participant answered the responsible-practices question by reference not to Japanese instruments at all but to its compliance with the EU General Data Protection Regulation (GDPR) and sectoral regulation in the US.

Participants report building their own governance programs in the following ways (note that participants were permitted to choose multiple options):



2.1.2 What the data tells us

Consulting Japan's voluntary AI guidance is not the same as being able to act on it. The survey therefore asked to what extent voluntary AI guidance, such as the AI Guidelines for Business (Version 1.2), helps participants to determine how to comply with "hard laws," such as the APPI, in their AI activities.

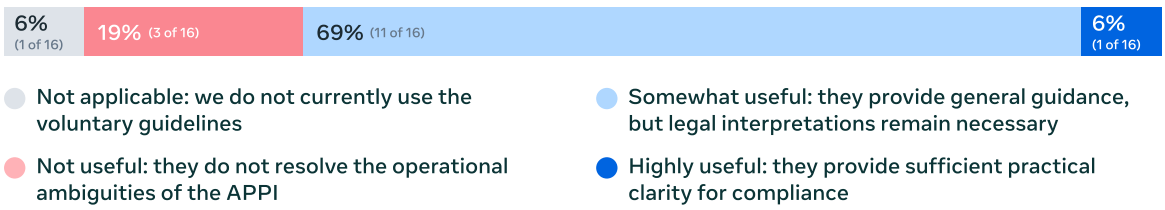


FIGURE 2.2
Usefulness of voluntary guidelines in navigating APPI compliance for AI

Set against the adoption figures, this is among the most striking results of the survey. It shows that voluntary AI guidance in Japan does appear to reach its audience – 11 of 16 (69%) read it (see Figure 2.2), and 3 of 16 (19%) treat it as their primary reference – but few (in this case, only a single participant (6%)) regard it as sufficient. 14 of 16 (88%) indicated that they either need legal advice on top of the guidance or find that the guidance does not resolve the questions they face.

Similarly, 4 of 16 (25%) identified aligning APPI compliance with the requirements of the AI Promotion Act and voluntary guidelines as a compliance challenge that they had encountered. 5 of 16 (31%) selected alignment of APPI obligations with the AI Promotion Act as a "ranked policy ask," and 5 of 16 (31%) also selected it as a "framework priority."

2.1.3 What challenges organizations are facing

The primary research identified three challenges that organizations are facing.

The first challenge that participants describe is one of translation and implementation

The guidance identifies risks but leaves each organization to identify mitigating practices to manage these risks.

- One participant put it precisely: Japan's voluntary AI guidance draws attention to the entry of personal and confidential information into AI systems, to the disclosure of important information when an organization integrates retrieval-augmented generation, and to privacy infringement. However, translating the guidance into its own operations requires it to prepare additional internal guidelines and checklists.
- A second participant suggested potential avenues to close the gap. These avenues include, for instance, asking for risk-based standards that organizations can implement in practice, standardized contracts and standard checklists, case studies, and sandboxes, to ensure safety while remaining agile.

An organization can accept a principle such as handling personal data appropriately in AI systems in full, and still not know whether a particular vendor arrangement is acceptable. Practical resources, such as model contract clauses, or checklists specifying what a reviewer must confirm before approving a use case, close that gap because it applies directly to the case in hand.

During the interviews, one participant described an uncertainty that can paralyze businesses while they

wait for guidance. The participant observed that the present period between promulgation of the APPI amendments and the new obligations taking effect raised the same challenge for its clients as did previous rounds of regulatory reform, including the last round of APPI amendments, and the introduction of the external transmission rules under the Telecommunications Business Act. In each case, clients' engagement with the new requirements was confined to their legal and compliance functions; the business functions that wanted to use the data were left unable to act; and most organizations settled into waiting for the outcome of the public comment process to resolve matters. The participant observed the same pattern now, with organizations again waiting for the public consultations, and on the subordinate rules and guidance that will follow from them to settle how the amended APPI applies in practice.

This account highlights that the period between promulgation of new amendments and their commencement is not dead time: it is when the subordinate rules and guidelines that determine how a reform applies are actually made. On this participant's evidence, organizations wait for these rules and guidelines before acting. A grace period is therefore most useful where the guidance arrives early enough within it to be absorbed, rather than shortly before the obligations bind.

A second challenge is that **guidance has not always kept pace with the technologies** that organizations are being asked to govern

In particular, **agentic AI** – systems that plan and carry out sequences of actions on their own – was cited as the clearest example in survey responses. Responses diverged on whether participants' existing compliance arrangements are suitable for such systems.

- 6 of 16 participants (38%) consider that the APPI's requirements do not, on their face, settle how they apply to systems that determine their own processing, and that working out the answer requires additional legal interpretation beyond what is stated in existing guidance.
- 5 of 16 participants (31%) specifically indicated that they are still working out how to give adequate notice and obtain consent for personal information processed by an AI agent.

- Taken together, 8 of 16 (50%) identified one or both of these as unresolved.
- On the contrary, only 2 of 16 (12%) consider that their existing APPI compliance measures already cover the behavior of autonomous agents. In interview, however, one of these two participants qualified that view: its assessment held for the systems within its own remit, and it acknowledged that use of AI tools elsewhere in the organization, outside its visibility, could change its answer.

No participant took both positions. Comparably placed organizations reading the same legislation have therefore reached opposite conclusions about whether what they already do is enough.

A further 5 of 16 (31%) report that they are not developing or deploying agentic AI at all. Most selected no other option in response to this question, so the survey alone does not establish why in those cases. However, one participant that selected this option nevertheless described, in a free-text answer, developing and operating agentic AI tools internally and testing a use case in which an agent autonomously consults internal databases to handle customer support enquiries. That participant added that the current APPI framework requires additional clarification as to how “acquisition” and “use” are to be interpreted where an agent accesses data on its own initiative.

The interview responses provided further explanations of the challenges faced by organizations.

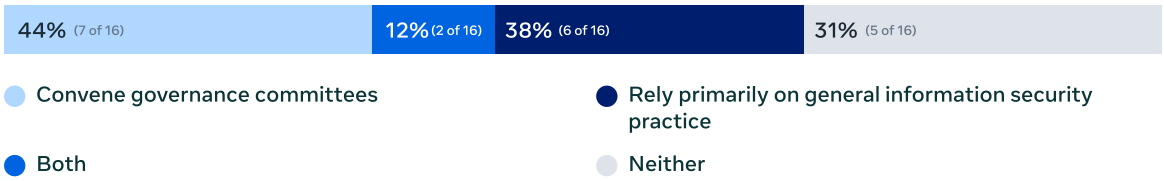
- One participant considered that the AI Guidelines for Business (Version 1.2) provide limited guidance on agentic AI, beyond defining the term.
- Two participants highlighted concerns about organizations’ ability to control AI agents. One, describing its own use of agents, said that complete control is difficult to achieve in practice. Agents may ignore an instruction not to take a particular action, or bypass explicit guardrails independently, rendering effective legal control highly challenging. Another, advising clients on AI adoption, reported growing concern that agents may operate outside the control organizations expect to retain over them, and that clients increasingly seek assurance that this risk can be managed.

The interview responses also suggested that the difficulty extends beyond agents. One participant observed that Japan's AI guidance is framed around generative AI, whereas the value organizations actually derive comes from combining generative models with conventional machine learning. The participant considered that physical AI will widen that gap further and suggested that guidance needs to rest on a broader account of what AI is.

2.1.4 What this means for the ecosystem

As discussed in **§2.1.1** above, Japan's AI governance framework currently relies on voluntary guidance to tell organizations what responsible AI practice looks like, which risks to watch for, and what good governance of an AI system actually means. Participants' experiences suggest that they read this guidance but must then instruct lawyers and convene internal committees to work out what the guidance requires of them. This results in uneven cost outcomes, as a large enterprise with a digital transformation division and an IT governance group would be able to absorb these follow-on costs far more readily than a startup, for instance.

The survey also shows that organizations differ in how they approach the work of AI governance. In the survey, Participants could select more than one practice, and the two most common were not mutually exclusive:



Some of those relying on general security practice build AI products or develop models themselves, and one processes biometric and health data. What separates the two groups is not scale or technical sophistication but whether the organization has concluded that AI needs governance distinct from information security governance.

One participant that relies primarily on general information security practice explained the reasoning directly. Its data security measures predate generative AI. It therefore regards AI as one tool among many rather than a distinct category of risk, and applies the same controls to data entered into an AI tool as to data entered into any other externally provided service. The participant added that its organization does not develop AI models to any significant extent, and that as a user of AI tools its position is not materially different from its position on any other software service.

This response suggests that significant implications arise from differences in companies' approaches to classifying the technology: an organization could treat AI as a new category requiring its own governance, or simply as a new instance of a category it already governs. If each organization builds its own operational interpretation alone, Japan may end up with a range of private-sector interpretations of what the APPI requires of AI systems, resulting in notably divergent practices towards AI governance that are difficult for a regulator to observe and track.

Asked what would make guidance usable, participants gave notably concrete answers.

- One participant called for **more accessible summaries**, drawing an analogy with the practice of government ministries. When legislation is made, the responsible ministry typically publishes the statutory text, a detailed slide deck, and a one-page summary to quickly communicate the essential points. The participant suggested that the same discipline be applied to guidance on AI and personal information.

- A second participant said the AI governance question that arises most often from its clients is whether there is a **spreadsheet checklist that can be used in day-to-day operations**. It contrasted two Japanese instruments, observing that the one closer to a checklist in format is the more useful or readily applicable of the two.
- A third proposed that the **guidance itself be delivered as an interactive tool rather than a document** – noting that the person completing a checklist in a smaller organization is typically an administrative or departmental representative rather than a specialist, and that a tool designed for the people actually building the product would be more likely to be used. The same participant suggested differentiating guidance by corporate function: setting out what a business development team should check, and at what point it should escalate to a management or compliance function.
- A fourth participant described a gap in existing guidance that the participant itself has been unable to address. The participant maintains assessment materials for AI readiness, but these are fragmented (e.g., across web services, applications, etc.). The participant is attempting to cover all of these within a single framework but found that nothing comparable is available in Japan. What it asked for was a comprehensive checklist or set of criteria capable of evaluating a complex AI service: one that incorporates risk assessment, states the standards that must be met before deployment, and could be put to use immediately. It added a requirement that bears on how such material would need to be built: AI is frequently not the main function of a service but one component within it, so the framework would need to scale to the level of AI involvement, identifying what is required in each case rather than treating every service alike.
  The same participant noted the limitation of standard material: a checklist allows an organization to demonstrate that it has met a minimum standard, but may not fit the circumstances of a particular company. It suggested case studies as a complement.
- A fifth participant made the same point from the perspective of an organization that uses AI rather than builds it. In its assessment, practical guidance directed specifically at organizations deploying AI as end users is lacking, and smaller organizations have limited resources with which to establish what the APPI requires of them. Its request was for a concise cross-border checklist for end-user organizations, and for sector-specific guidance on AI utilization in its field.

Opportunities for the PPC and other policymakers

OPPORTUNITY 1

For METI and MIC:

> ## Supply practical resources in the next iteration of the AI Guidelines for Business

To supplement METI's existing Contract Guidelines on Utilization of AI and Data, organizations would benefit from access to additional practical resources that they can apply directly in their operations. These could include worked examples, model contract clauses, standard review checklists, and sector annexes. In particular, an explicit mapping between guidance in the AI Guidelines for Business (Version 1.2) and corresponding binding obligations under the APPI would allow organizations to address both frameworks in a single internal review. The AI Guidelines for Business would also benefit from greater guidance on agentic AI.

OPPORTUNITY 2

For the Cabinet, AI Strategy Headquarters, METI, MIC, and the PPC:

> ## Establish a sandbox to test innovative ideas, put in place structured dialogue pathways for regulators and industry to discuss novel or complex compliance queries, and publish case studies of compliant AI deployment in Japan

These three measures address the same gap from different directions: organizations cannot establish in advance whether a novel approach is acceptable, and there is no route by which they can find out other than commissioning legal advice on a question the guidance does not answer.

A sandbox to test innovative ideas

3 of 16 participants (19%) selected sandboxes for testing compliance approaches before full rollout, and one set out what such an environment would need to permit: organizations within the same corporate group sharing data between entities for a single AI development project without each transfer requiring separate treatment, and unrelated organizations pooling data for a defined development purpose.

OPPORTUNITY 2 CONTINUED

The new statistical creation exemption, added by the APPI Amendment Bill, appears to provide for much of this, permitting the provision of personal information to a third party for statistical purposes without consent where both parties publish prescribed information about the activity and a written agreement records that the provision is made on that basis. What would remain useful is a controlled environment in which organizations can establish, before committing, whether a proposed collaboration falls within the exemption: what must be published, what the agreement must record, and how the purpose limitation and onward transfer restrictions apply to a multi-party project.



Structured dialogue pathways

Participants consistently described a mismatch between the questions they have and the answers available. To address this, participants proposed measures that supplement guidance rather than adding to it: a channel through which a specific question can be put, and an answer relied upon.

- One participant asked for the expansion of advance consultation with the PPC, so that an organization could consider appropriate governance measures before implementing them, rather than after a problem has emerged. The same participant noted that most smaller organizations cannot afford to engage a large consultancy or law firm.
- A second participant ranked industry roundtables with regulators second among its policy asks and described what it would want from such a forum: that the organizations supplying AI systems disclose, in transparent form, what data they used, what they did with it, and how trust in the resulting system is assured.
- A third participant pointed to the Open Loop Japan project itself as a model for collaboration between industry, government, and academia. The participant proposed that the government lead this forum, emphasizing that cross-sectoral collaboration is essential for its success. The rationale for this initiative is specific: while public discourse often gravitates toward restrictive regulation, the focus should instead remain on expanding the potential for innovation. Given that inaction is frequently the path of least resistance, there is a recurring psychological tendency to default to conservative restrictions. The participant accordingly sought a multi-stakeholder environment, encompassing the state and academic institutions, capable of establishing a more constructive and enabling direction for the ecosystem.

What these accounts have in common is that the questions organizations cannot answer are novel and specific: whether a particular vendor arrangement is acceptable, whether a proposed identifier solution can be used in Japan, whether a given level of de-identification suffices for a given recipient. Guidance addresses classes of questions. A dialogue pathway addresses instances, and can do so before an organization has committed to a course of action.

OPPORTUNITY 2 CONTINUED

Published case studies

2 of 16 participants (12%) selected case studies or benchmarks showing compliant AI deployment in Japan. Asked what would make standard material usable, one participant observed that a checklist allows an organization to demonstrate it has met a minimum standard but may not fit the circumstances of a particular company and suggested case studies as the complement. These are modest numbers, but the request is for something no organization can produce alone.

2.2 Finding 2

Organizations have difficulty establishing how far they must de-identify personal information to use it in their AI activities

The cohort's highest-”ranked policy ask”, by weighted preference, is a practical, objective, risk-based standard for what counts as sufficient anonymization under the APPI. The evidence suggests that participants are willing to take steps to reduce the identifiability of personal information. However, organizations face difficulties in establishing in advance how far they must de-identify data to meet the APPI’s standards for anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報). Faced with this uncertainty, some organizations have stopped trying to rely on these provisions and in the process, may forgo the flexibility that the APPI’s anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報) regimes were designed to provide.

2.2.1 What organizations are already doing

Some participants take steps to reduce the identifiability of personal information used in AI- related activities. Just under a third of participants – 5 of 16 (31%) – use only anonymized or aggregated data for AI purposes.

Where personal information is provided to a third party engaged to carry out AI development, participants apply a mix of measures:

- **Encryption**, which protects the data but does not change its status under the APPI;
- **De-identification** that does not meet a statutory standard; and
- **Conversion into anonymously processed information (匿名加工情報)**, which is the only option of the three that formally takes the data outside the APPI's scope.

Participants sometimes apply these in combination and not always consistently.

During the interviews, one participant set out how it approaches this. Depending on the circumstances it may apply any of those techniques, but it does not rely on any of them to alter the data's status. Its working assumption is that personal data remains personal data however far it is processed, and it therefore proceeds by treating all such information as though it is subject to the full suite of obligations under the APPI whenever it provides the data to a third party.

A second participant described a comparable position. It told us that its organization has pseudonymization and anonymization processes in place and that it applies the measures it considers necessary on a project-by-project basis. Nevertheless, it also continues to treat that data as though it is subject to the full suite of obligations under the APPI. The participant was explicit that this is deliberate over-caution, rather than a considered reading of what the law requires.

The reasoning behind these two responses to an indeterminable standard is understandable. Its practical effect, however, is to forgo the flexibility and benefits that the APPI's anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報) regimes were designed to provide.

Separately, privacy-enhancing technologies (PETs) remain a low priority for this cohort. Across the policy-related questions in the survey, only 1 of 16 participants (6%) called for greater regulatory guidance on PETs, and 3 of 16 (19%) called for incentives to adopt them. This is notable given the nature of the difficulty participants describe. PETs are precisely the class of technique that would allow an organization to demonstrate what a given processing operation achieves, rather than asserting it. That they attract so little attention suggests the constraint is not a shortage of available techniques but **the absence of any external standard against which the output of those techniques could be measured.**

2.2.2 What the data tells us

The survey asked participants to rank the policy changes and support mechanisms that would most benefit their AI operations. All 16 completed the ranking, and we scored the results by weighted preference.

Practical, objective, risk-based standards for "sufficient anonymization" under the APPI ranked first by weighted preference. 10 of 16 (62%) selected it, and 4 of 16 (25%) ranked it first. No other option attracted more selections or a higher weighted score.

Separately, when asked what compliance challenges they had encountered under the APPI, 8 of 16 (50%) participants identified challenges in determining how far data must be de-identified to fall outside the APPI's scope. 8 of 16 (50%) also selected a flexible, practical, context-based and risk-proportionate definition of anonymization as a “framework priority.”

2.2.3 What challenges organizations are facing

The APPI already contains processing standards for anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報). However, participants report several operational uncertainties, such as the required extent of processing, applicable datasets, permitted recipients, and allowed purposes.

The first is the substantive threshold for anonymization

One participant reports that it encrypts, anonymizes, or pseudonymizes the data before providing the data to a third party engaged to carry out AI development. However, in that participant's view, how far this processing must go – and, assuming it is required at all, why – has never been clearly articulated. The participant regards this as a point that warrants considerably more specific attention.

The second is the standard's character

A second participant asks for clarification, in a form that does not impede AI use, of how far it must de-identify data before it can use it under the APPI, on a practical and risk-based footing – expressly so that organizations do not become excessively conservative and can instead promote data use under appropriate governance.

The third is assurance

A well-drafted standard remains of little value if no external party can confirm it has met it. A third participant spoke directly to that concern, calling for software that can demonstrate that an anonymization or pseudonymization product performs as intended, and for clarity about which public or third-party accreditation bodies stand behind such tooling, with criteria and check items published in advance.

The effects of these three difficulties are cumulative. An organization that cannot determine the threshold for anonymization, cannot predict how it will be assessed, and cannot obtain external confirmation that it has been met is left to set its own standard and carry the consequences of being wrong.

2.2.4 What this means for the ecosystem

The request here is for determinable criteria, not for lower standards, and one participant said plainly that the risk it wants addressed is organizations becoming too conservative because the threshold for anonymization is unclear.

Participants respond to the uncertainty in different ways, and the two most visible responses point in opposite directions while producing the same effect. Some restrict AI work to anonymized or aggregated data, staying outside the APPI's scope by declining to bring potentially richer data into it. Others assume that data never leaves the APPI's scope at all, and treat everything as personal data throughout. Both are rational answers to a threshold for anonymization that cannot be established in advance, and both give up something: the first forgoes the analytical value of the data, the second forgoes the flexibility the statutory regimes were designed to offer.

What neither approach produces is a shared understanding of where the line falls. Where the threshold for anonymization cannot be determined before the fact and cannot be validated after it, organizations facing the same rule reach materially different positions on the same data, and an individual's protection comes to depend on which position the organization handling their data has adopted.

Recommendations for the PPC

RECOMMENDATION 1

Publish practical, risk-based criteria for how far an organization must de-identify personal information before they can use it in AI-related activities

Thresholds should be proportionate to risk – avoiding over-conservative defaults that shrink usable data – and state the legal effect of being met, so organizations can rely on them in advance (e.g. a safe harbor).

To be most valuable, organizations would need to be able to apply those criteria before the fact. As seen above, the key challenge as described by participating companies is that they cannot establish in advance whether they have met the relevant threshold.

Three elements would make such guidance operational:

Tie data thresholds to the risk level of the use case

A single standard for internal analysis, sharing data with partners, and public release will inevitably be too strict for internal use or too lenient for published data. Differentiating between these different use cases allows the standard to be calibrated to the risks that each may present. Participants asked for standards to be practical, risk-based, and achievable. In practice, this would minimally require that the standard should be based on factors that organizations can easily assess for themselves (e.g., who will receive the data, what other data the recipient holds, and whether the two could be matched to identify individuals) and should focus on the risk of re-identification. Further, in terms of format, criteria could take a form that organizations can check against easily, rather than a high level principle that organizations are required to interpret and apply.

Provide accessible guidance on the legal effect of the APPI's anonymously processed information (匿名加工情報) and pseudonymously processed information (仮名加工情報) regimes

Explain what an organization gains by converting data into anonymously processed information (匿名加工情報), what it gains from pseudonymously processed information (仮名加工情報), and how the legal position of that data changes when it is processed without meeting either standard. This would allow organizations to choose deliberately between the two regimes, rather than defaulting to the most conservative available reading.

RECOMMENDATION 1 CONTINUED

State the legal effect of meeting the criteria

Participants invest in de-identification without knowing what legal position that investment secures. Guidance should say what changes once a standard is met: whether the data may then be provided to a third party without consent, whether it may be used for a purpose beyond the one specified at collection, and what obligations if any continue to apply. Without that, an organization has no way to weigh the cost of processing data further against any benefit from doing so.

Opportunities for other policymakers

OPPORTUNITY 3

For METI, MIC, the PPC, and standards bodies:

Establish external assurance for de-identification tooling

Publishing criteria alone leaves organizations to certify their own compliance. One participant highlighted the need for external assurance: it can find neither tooling that demonstrably performs anonymization and pseudonymization correctly, nor a public or private body that assesses such tooling against published criteria.

A safe harbor would address this directly: an organization that applies published criteria in good faith would be protected, even if a later assessment reaches a different conclusion. This would give organizations a basis for confidence beyond their own judgment.

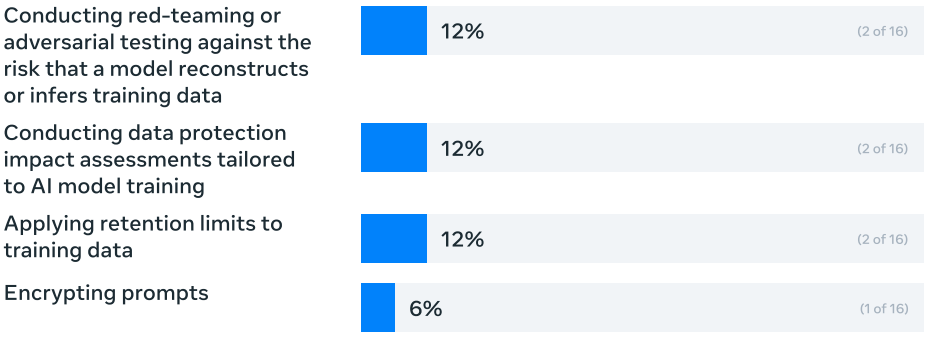
De-identification is a fast-moving field in which both the techniques for de-identifying data and the methods for reversing it continue to advance, and any fixed assessment of a tool may cease to reflect the state of the art well before it is revisited. Any assurance route should therefore be designed to be revised and should recognize the international certifications many organizations already hold (such as ISO/IEC or SOC standards) rather than duplicating them.

2.3 Finding 3
The classification of AI outputs and models under the APPI remains unresolved

Among the most frequently cited compliance challenges in the survey is the risk that generative AI outputs reconstruct or infer personal information. Beneath this, the status of trained AI models and their outputs under the APPI remains unclear. Two participants independently identify this as an unresolved issue that impacts their operations. One reports proceeding with deployment notwithstanding this uncertainty.

2.3.1 What organizations are already doing

Generally, participants manage "**output related risks**" – the risks that generative AI outputs could reconstruct or infer personal information – using a range of tools and techniques:



Beyond these, the broader survey results suggest that organizations may be managing these risks by restriction: confining generative AI to internal environments, to non-sensitive workloads, or only using anonymized information in connection with AI workloads.

2.3.2 What the data tells us

8 of 16 participants (50%) cited output-related risk as an uncertainty and compliance challenge that they encounter under the APPI in their AI-related activities. This places output-related risks in joint first position among such challenges, tied with anonymization (see **Finding 2**), and ahead of cross-border data transfers at 7 of 16 (44%) (see **Finding 4**).

2.3.3 What challenges organizations are facing

Two participants set out the problem in detail but approached it from different directions.

> One participant explains that it fine-tunes particular models to use generative AI internally, and has deployed the result across the organization

However, the participant expressed uncertainty as to **whether outputs returned by the model constitute personal information**. The participant explained that it has continued to use the system without settling that question.

In the interview, the participant expanded on this: there is no ruling, guideline, Q&A or evaluation criterion from the PPC addressing the status of information an AI system returns. As a result, it could not determine whether its position was permissible or not. It also said that if output were interpreted as personal data held by the organization, providing that output to a

third party would itself require the individual's consent, which would restrict how the output could be used.

Asked how the risk is managed internally, the participant shared that it has not issued detailed guidance on what to do if privacy-infringing information appears in an output. It described this as something it hopes to address more thoroughly but has not yet been able to.

Another participant addresses the legal classification of a trained AI model under the APPI

This turns on the APPI's unique distinction between "personal information" (個人情報), "personal data" (個人データ), and "retained personal data" (保有個人データ)

Under Article 2(1) of the APPI, "**personal information**" (個人情報]) is the broader, foundational category encompassing any information relating to a living individual that can identify a specific person. Obligations for "personal information" focus on transparency and scope, requiring businesses to clearly define their purpose for collection and restrict usage to that stated purpose.

Conversely, "**personal data**" (個人データ) is a subset of "personal information." According to Article 16(3) of the APPI – which builds upon the definition of a "personal information database" (個人情報データベース等) in Article 16(1), "personal data" refers exclusively to personal information that has been compiled and systematically organized into a searchable database. Obligations for "personal data" are more stringent. Businesses must implement robust security measures, obtain consent for transfers to third parties, and report data breaches.

If a business has the authority to disclose, correct, add to, delete, or suspend the use of "personal data" (個人データ), that personal data is classified as "**retained personal data**" (保有個人データ) under Article 16(4) of the APPI. At this level, the business must accommodate data subject rights, meaning individuals can legally request access to, or correction or deletion of the data.

The participant reads a notice from the PPC on the use of generative AI services (released in June 2023)

as implying that a large language model (**LLM**) falls within the APPI's definition of a "personal information database." The participant noted that the PPC's official interpretation is that search engines fall outside of this definition because they are not arranged systematically for retrieving personal information. Since a LLM is likewise not arranged for that purpose, the participant reasons that LLMs too should fall outside the definition. However, the participant has observed that practice following the alert treats such LLMs as though they qualify as "personal information databases." Interpretations of whether LLMs constitute personal data differ internationally, and the participant considers that Japan needs clear guidance on this.

The participant highlighted that if LLMs do qualify as "personal information databases," it is unclear how organizations should respond to requests under the APPI to disclose, correct, or suspend their use as "retained personal information." Because that interpretation remains unsettled, the participant cannot set a direction for business design.

In interview, the participant did not consider it feasible for a deployer to give effect to disclosure, correction and suspension requests as the deployer simply does not hold the model in a form that would allow it. Therefore, the obligations attaching to disclosure and correction would in practice have to fall on the developer of the underlying model.

The participant also considered that because the APPI's distinction between personal information, personal data, and retained personal data is unique to Japan's APPI, a developer of a foundation model, typically operating globally, is unlikely to construct compliance mechanisms for a distinction that exists in one jurisdiction. In the participant's assessment, the obligation would therefore attach to a party with little incentive to discharge it, and the deployer facing the individual would be unable to give effect.

These are two independent participants, in unrelated sectors, describing the same gap. The second account also shows why this is more than a classification question. If a trained model is a personal information database containing retained personal data, then the APPI's disclosure, correction and suspension rights attach to it – and no participant identified a way to give effect to a deletion request against model weights. The participants perceive that the rights and the technology, at present, seemingly do not align.

2.3.4 What this means for the ecosystem

The consequence of this unresolved classification of AI models and outputs in relation to APPI is that organizations forgo opportunities to employ AI in situations where they otherwise would have done so.

Participants describe this directly. One confines generative AI to work outside its sector's regulated quality framework, pending sectoral guidance on how that framework applies. Another has an internal use case involving employee data that has not proceeded. 5 of 16 participants (31%) restrict AI work to anonymized or aggregated data, and 5 of 16 (31%) process no personal information in connection with AI at all. One participant described the barrier inside its own organization as psychological rather than legal: where staff interpret personal information broadly and rule out anything involving AI on the assumption that it would be too risky, useful applications are abandoned before anyone assesses them.

The unresolved status of models and their outputs contributes directly to this. An organization that cannot determine whether what its model returns constitutes personal information cannot determine what obligations attach to it, and the conservative course is to narrow the use case or not to build it. Two participants did continue on an unresolved footing, which carries its own risks – for the organization, and for individuals whose position depends on which interpretation was adopted. However, the more common response, on this evidence, appears to be restraint.

There is a reasonable argument that the underlying question has a narrower answer than current practice assumes. One participant set it out: the official interpretation is that search engines do not constitute a “personal information database” (個人情報データベース等) under the APPI because they are not systematically organized for the retrieval of personal information, and a trained model is not so organized either. On that reasoning, a model would fall outside the definition. What matters, then, is whether a particular output identifies an individual, rather than what the model contains. The same participant noted that practice following the PPC’s June 2023 alert has tended the other way, and that interpretations differ internationally. Resolving this would allow organizations to direct their effort at the point where the risk actually arises, which is the output.

Recommendations for the PPC

RECOMMENDATION 2

> **Consider issuing guidance clarifying that trained models and their outputs are generally not personal data absent identifiability, and consider issuing guidance on practical measures to minimize output inference risk**

The new statistical creation exemption applies only where the output of the activity is not itself information relating to an individual, and the basis on which AI development is treated as falling within that definition is that a trained model represents statistical patterns rather than individual records. Confirming that a trained model is not, of itself, information about an individual would therefore align the treatment of models with the premise on which the exemption already rests.

Confirm the legal classification of LLMs under the APPI

Guidance confirming that a trained model does not, of itself, constitute a personal information database under the APPI – and that its outputs are personal data only where they identify an individual – would resolve the question participants were least able to settle. The reasoning is available in the PPC's own existing interpretation: search engines fall outside the definition because they are not systematically organized for the retrieval of personal information, and a trained model is not so organized. Confirming this would align the treatment of models with the treatment of an analogous technology, and would not reduce protection for individuals, because the obligations that matter would continue to apply wherever an output identifies someone.

Guidance should also address where in a development and deployment chain any obligation is intended to sit. One participant observed that if a model does constitute a personal information database, the obligations of disclosure and correction should fall on its developer rather than on the organization deploying it as the latter cannot technically satisfy these requests.

Address intermediate outputs

The carve-out in the statistical creation exemption – that the output must not be information relating to an individual – is straightforward for a trained model. It is less so for embeddings, feature vectors, model activations and similar intermediate artefacts, which represent statistical relationships but may in principle be linked back to particular training data. An organization relying on the exemption needs to know whether producing or retaining such artefacts is consistent with that requirement. This is the same question participants raised about model outputs, arising at a different point in the pipeline, and it will determine how far the exemption can be relied on in practice.

RECOMMENDATION 2 CONTINUED

Provide practical guidance on measures to minimize output inference risk

With the status question settled, the substantive work is at the output, and this is where organizations are least equipped. 8 of the 16 participants (50%) identified the risk that generative AI outputs could reconstruct or infer personal information – the most-cited challenge in our survey, along with anonymization – while far fewer have measures directed at it: 2 of 16 (12%) conduct red-teaming or adversarial testing against models leaking or reconstructing training data, and 2 of 16 (12%) apply retention limits to training data.

The gap participants describe is specific. One told us its policy directs staff not to enter personal information as input, but that it has issued no guidance on what to do if privacy-infringing information appears in an output – something it considered it should address and had not been able to. Another reported that anything in an output that might resemble personal information has to be examined. A third has issued awareness-level warnings, its concern being that data may be incorporated and used internally without the organization being fully aware of it.

Guidance would be most useful if it identified the measures organizations can adopt and indicated what a proportionate level of each looks like. On this cohort's evidence, the relevant measures include:

- **Input controls** – restricting what may be entered into a system, and stating the limits of that approach where a model is fine-tuned on the organization's own data;
- **Pre-deployment testing for memorization and reconstruction** – with an indication of what constitutes proportionate testing, since 2 of 16 participants (12%) already do this and there is no stated expectation against which to calibrate;
- **Retention and deletion practices for fine-tuning data**, and their effect on what a deployed model can reproduce; and
- **Output handling** – what an organization should do where personal information does appear in an output, including whether provision to a third party has occurred; and
- **Avoiding re-training on inputs** – one participant told us it does not enter data into generative AI systems that will be re-trained on it, and avoids such systems. Whether that is necessary, or whether contractual assurance suffices, is a question guidance could settle.

2.4 Finding 4

Cross-border rules prevent participants from using AI systems hosted outside of Japan, and cloud rules cause uncertainty and hesitancy

Cross-border data transfer requirements are a key area where participants report genuine impact on their decisions to develop or deploy AI. Just under half of the cohort reports uncertainty about which overseas providers and jurisdictions satisfy the APPI's equivalent protection and equivalent measures requirements, and difficulty in applying these to specific vendor arrangements. None of the 16 participants use the Global Cross-Border Privacy Rules (CBPR) framework, but 11 of 16 (69%) identify a single clarification that would change that: PPC confirmation that a data exporter in Japan may transfer personal information to a recipient holding Global CBPR certification without needing any other legal basis or complying with any other legal process or requirement.

2.4.1 What organizations are already doing

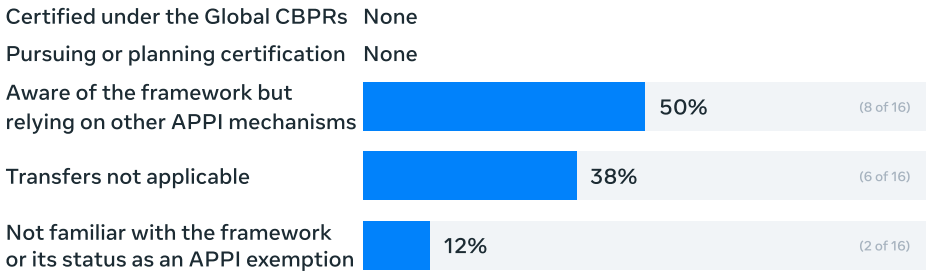
4 of 16 participants (25%) transfer data to overseas cloud infrastructure for AI processing, and 1 of 16 (6%) handles personal information specifically at the cross-border transfer stage of its AI workflow.

Participants perceive difficulties in identifying an appropriate legal mechanism for transferring personal information outside Japan.

Under Article 28 of the APPI, these mechanisms include:

- "**Adequacy decisions**" (十分性認定), which allow organizations to transfer personal information to jurisdictions that the PPC has designated as offering equivalent protections to the APPI,[5] without any further restrictions;
- In the absence of such a designation, establishing that the recipient has implemented **equivalent measures** (相当措置) to protect the personal information post-transfer or obtaining **informed consent**.

Participants were also asked to share their positions on the Global CBPR, an international certification framework for organizations transferring personal information across borders that could satisfy the "equivalent measures" requirement:



2.4.2 What the data tells us

7 of 16 (44%) identified cross-border clarity as a compliance challenge they had encountered - placing this in the top three challenges faced overall (see above).

In terms of specific compliance challenges faced in this context, participants' responses are outlined in Figure 2.4 below.

[5] Currently, the European Economic Area, and the United Kingdom.

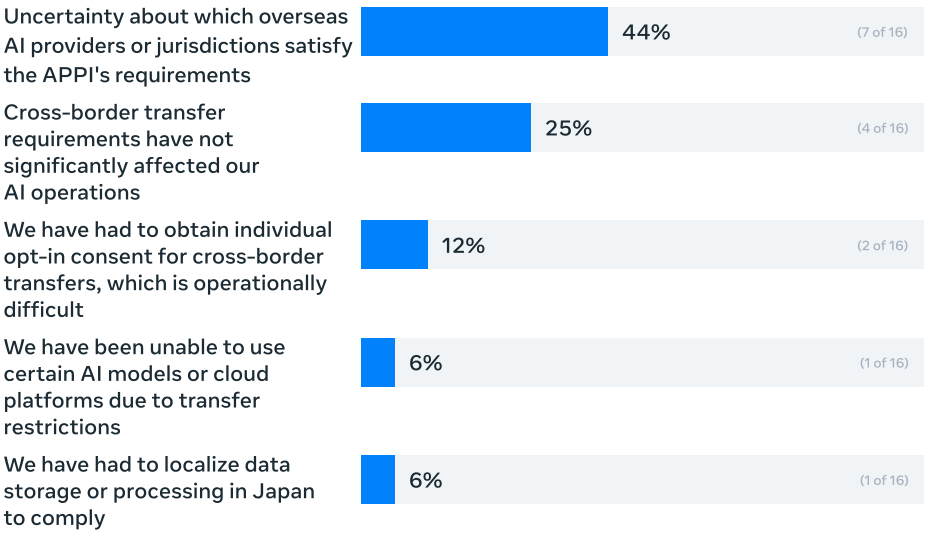


FIGURE 2.4
Effect of the APPI's cross-border transfer requirements on AI-related operations (n=16, multiple selection)

Asked whether the uncertainty concerns the recognition of jurisdictions as adequate or the measures taken by individual providers, one participant said both. Approaches to data protection differ by country or region, and reconciling those differences when transferring data globally is demanding; separately, the security measures implemented by each particular recipient vary, and those variations have to be managed as well. The participant described these as two distinct challenges.

In questions on policy preferences, participants called for more straightforward transfer rules for AI.

- 6 of 16 participants (38%) selected a streamlined interoperable framework for cross-border data transfers in AI contexts as a “ranked policy ask”, with 2 of 16 (12%) ranking this first out of fifteen options in priority.
- 8 of 16 (50%) selected streamlined transfer rules enabling global AI collaboration as a “framework priority.”

2.4.3 What challenges organizations are facing

One participant’s account identifies three separable difficulties:

Factual

Working out where personal information actually goes once it enters a vendor's service – which sub-processors handle it, in which countries they operate, and whether the service keeps it in logs or uses it beyond the immediate request – as distinct from what the service description states.

Contractual and technical

Reflecting those flows in contractual terms and technical and security measures that satisfy the APPI's equivalent-measures requirement.

Commercial

An organization's ability to obtain the contractual terms it needs depends on its bargaining position with the provider, and the largest AI providers rarely negotiate.

It should also be noted that this participant is among one of the larger organizations in the cohort, which suggests the bargaining constraint is not limited to smaller organizations.

A second participant considers that:

The APPI's transfer rules themselves do not align with technical realities

Achieving a single purpose of use now routinely requires personal information to move across and be processed by several business operators. In this participant's view, the difficulty arises from the architecture of the APPI: because the APPI regulates each business operator separately, each business operator assesses its own position independently, and no assessment is made of the processing chain as a whole, even where several operators contribute to a single purpose. The participant contrasted this with the GDPR, which in its account is oriented around the processing operation rather than the organization, so that obligations track what is being done with the data rather than which entity holds it at a given moment. In this participant's account, protection across the chain is not coordinated under the APPI's approach, since each operator reasons only about its own segment. Further, obligations attach most heavily at organizational boundaries rather than in proportion to risk, so the same data used for the same purpose attracts fresh requirements at each handover.

A third participant reported that:

The first question Japanese organizations always ask about a service is whether the data center is located in Japan

In a similar vein, a fourth participant described a specific decision to localize data in Japan.

- The participant has recorded certain business conversations and needed those recordings transcribed. The transcription services performing best at the time were provided by organizations based in the United States. Because the recordings were of business conversations, they necessarily contained the names of participants on both sides, and masking every occurrence of a name before sending material overseas was not practical. The participant used a domestically provided service with lower performance instead.

- The same participant described its general position as follows: unless there is a compelling reason otherwise, it uses cloud services in Japan only. Its reasoning was that transfer to an overseas service provider would require the individual's consent even where the arrangement is a service contract, and obtaining that consent would be difficult in practice. It has effectively ruled out services it cannot access domestically.

The foregoing discussion assumes that the APPI's transfer rules apply at all. This was also cited as an area of uncertainty.

A fifth participant questioned:

Whether use of an AI service could fall under the so-called "cloud exception" in Section 7-53 of the PPC's Q&A guidance

Under this provision, the PPC does not consider certain transfers of personal information to cloud service providers as a provision of personal information to a third party (which generally requires individual consent under the APPI) if the provider does not handle the personal information for its own purposes. The participant described the exception as unique to Japan and considered it challenging to interpret. If the cloud exception applies to use of AI services, then a substantial volume of ordinary AI activity – inference API calls, fine-tuning on a vendor platform, retrieval over an external index, log retention – may not engage the APPI's cross-border rules at all. Some of the uncertainty at **Figure 2.4** may therefore also include uncertainty about whether the transfer rules apply, rather than about how to satisfy them.

A sixth participant identified a related uncertainty that arises earlier in the analysis:

Whether sending data to an AI service provided from overseas amounts to a cross-border transfer at all

Its difficulty concerns the content of prompts. Where an employee submits a summary of customer information to a LLM API operated by an overseas provider, the participant could not determine how far that constitutes a provision of personal information. It asked for practical guidance addressed to organizations using AI as end users, in the form of decision flowcharts and concrete questions and answers – naming, as the questions that recur in its own operations, how far data must be de-identified before it ceases to be personal information, and whether submitting data to an API constitutes provision to a third party.

A seventh participant suggested that:

AI adoption has caused industry practice to change

Observing that many organizations that, in the software-as-a-service (SaaS) context, previously paid close attention to whether personal information was transferred to overseas servers now approach AI as though that experience had not happened.

Against this, the Global CBPR result stands out. As mentioned in **§2.4.2** above, a majority of participants called for simpler transfer rules for AI, and the Global CBPRs – a certification framework designed to deliver exactly that – already exists. Yet, none of the 16 participants hold, or are seeking, Global CBPR certification.

Uncertainty about what the certification achieves may explain the gap: 11 of 16 participants (69%) called on the PPC to confirm that a data exporter in Japan may transfer personal information to a recipient holding Global CBPR certification without needing any other legal basis or complying with any other legal process or requirement. These results are noteworthy considering that the PPC clarified the position of the Global CBPR System under the **APPI through amendments to its APPI Guidelines (Transfers to Third Parties Located Overseas)** in December 2025. This suggests that the policy clarification has not been accessible to a broader range of organizations, making this a question of visibility rather than law.

However, participants also shared views on why uptake of CBPR certification has been limited in Japan.

- One participant highlighted the lack of a clear value proposition for CBPR certification. Without a defined benefit, it said, an organization has no motivation to pursue certification. It identified two things that would constitute such a benefit: exemption from specific requirements, or an alteration in how responsibility is determined and allocated in the event of an incident; and broader recognition of the framework among governments and users, such that certified organizations were understood to be safer. On the second point the participant observed that if enough people shared that understanding, certification would become worth obtaining for that reason alone.

- A second participant located the obstacle in capacity rather than in legal effect. Obtaining a certification is one matter; maintaining it requires an internal system the organization has not established, and it does not currently see a return that would justify building one – it does not provide AI services directly to consumers, so the value of certification depends on how its business develops. What would change that calculation, in its view, is counterparty pressure: if the organizations it does business with held such certifications, obtaining one would become close to unavoidable.

2.4.4 What this means for the ecosystem

Taken together, these answers describe a difficulty of application rather than of knowledge. The rules and the designations are published. What consumes participants' capacity is tracing where the data goes, drafting terms to align with these realities, and negotiating them with the provider. Participants report having to conduct the same assessment for each new vendor and service. Organizations best placed to absorb these obligations are the ones with dedicated privacy functions, which mirrors the unevenness that **Finding 1** identified in governance practice.

It is also noteworthy that participants that have taken divergent approaches to complying with Article 28 of the APPI and have reported very different effects on their operations: 4 of 16 (25%) report no significant effect; 2 of 16 (12%) now obtain individual consent for each transfer; and 1 of 16 (6%) has moved processing into Japan. Where the contractual terms an organization can obtain depend on its bargaining position, how well data is protected may follow the exporter's commercial strength rather than the sensitivity of the data.

On the Global CBPRs in particular, an existing certification mechanism that no participant has is not one that nobody wants but rather, one whose value and legal status could benefit from greater clarity and awareness.

Recommendations for the PPC

RECOMMENDATION 3

Make the cross-border transfer mechanisms (e.g. model clauses, Global CBPR certification) easier to adopt and determinable in advance

We welcome the PPC's commitment in its 2026 International Strategy to develop model contractual clauses for transfers of personal information out of Japan.

Participants describe two distinct problems. The first is that the mechanisms available to them are harder to adopt than the alternative of avoiding transfers altogether. The second is that even where a mechanism is adopted, its legal effect is not determinable before the fact. Both need addressing and require different measures.

Making the mechanisms easier to adopt

1. **Publicize the policy clarification regarding the effect of Global CBPR certification**
   Amendments to the PPC's guidelines made in December 2025 have already clarified the legal effect of Global CBPR certification. However, notwithstanding this policy clarification, a clear statement that organizations can rely on the Global CBPR to transfer data across borders without engaging another transfer mechanism or legal requirement under the APPI is one of the most requested policy interventions in the survey: 11 of 16 participants (69%) specifically asked for this. That figure is itself evidence about visibility rather than about the substance of the rules. A clarification already made, which the organizations it was made for are still asking for, has not reached its intended audience. Only 2 of 16 participants (12%) told us they were unfamiliar with the framework, yet none is certified or currently pursuing certification. This is a communication problem rather than a legal one, and the remedy is correspondingly straightforward.

2. **Provide incentives for adoption of Global CBPR certification**
   Clarity on the legal effect of Global CBPR certification is important but is only the starting point for encouraging organizations to seek Global CBPR certification. As set out at **§2.4.3**, participants identified two distinct obstacles to pursuing certification:

   (1) a perceived absence of benefit from holding certification, and
   (2) the ongoing capacity needed to maintain one where the return is not yet visible.

   Neither participant objected to the framework itself. The difficulty is that certification does not yet secure anything an organization can point to justify pursuing certification, such as an exemption from other requirements, a change in how responsibility falls if something goes wrong, or broader recognition in the market.

RECOMMENDATION 3 CONTINUED

Streamlining the certification process and extending the validity period would reduce the maintenance burden participants describe. A safe harbour for processing covered by a certification, where the organization acted in good faith and with due diligence, would also provide a clear value proposition. Additionally, many organizations already hold international certifications (e.g., ISO/IEC 27701) covering adjacent ground. Publishing a mapping of which Global CBPR requirements those certifications already satisfy, and permitting existing audit evidence to be relied on where it covers the same requirements, would lower the cost of certification.

Uptake may also be self-reinforcing. One participant observed that if its counterparties held certification, obtaining one would become close to unavoidable – which suggests that adoption among a small number of large purchasers or providers could do more to drive uptake than any incentive offered directly to individual organizations.

3. **Make provider adoption of the (forthcoming) model clauses observable**
   Model clauses relieve the constraints participants describe only if providers accept them. One participant reports that the terms an organization can obtain depend on its bargaining position, and that the largest AI providers rarely negotiate. Publishing the clauses in English alongside Japanese, and maintaining a public record of providers that have adopted them unamended, would convert an exporter's task from opening a negotiation into checking a list.



Making the mechanisms determinable in advance

1. **State the legal effect of adopting the model clauses**
   Whether using them discharges the equivalent-measures requirement, and what if anything an exporter must establish in addition. This determines whether an organization can rely on them at the design stage or must obtain advice in each case.

2. **Draft the (forthcoming) model clauses with AI processing patterns in view**
   To address the arrangements participants describe, the clauses would need to cover sub-processor chains and the patterns of processing that characterize AI workflows, including inference API calls, fine-tuning on a vendor platform, retrieval over an external index, and log retention.

3. **Publish model criteria for provider-level assessment of equivalent measures**
   Where an organization does not rely on a certification, guidance on what it should establish before contracting with overseas providers: if the exporter does not have Global CBPR certification, what documentation suffices, how far a chain of sub-processors must be traced, and what a suitable answer looks like, would be helpful, especially for organizations without dedicated privacy functions.

RECOMMENDATION 3 CONTINUED

4. **Restate the cloud exception against the processing patterns AI actually produces,** such as inference API calls, fine-tuning on a vendor platform, retrieval over an external index, and log retention. This is among the highest-value clarifications available, because it governs a large share of routine AI processing, and determines whether the transfer rules are engaged at all. At least one participant reports being unable to interpret it.

Opportunities for other policymakers

OPPORTUNITY 4

For METI and MIC:

Align the Contract Guidelines on Utilization of AI and Data with the PPC's (forthcoming) model clauses

We welcome METI's Contract Guidelines on Utilization of AI and Data, which give parties a reference point for allocating rights in data and in trained models.

The Guidelines address contracting between commercial parties. The PPC's model clauses will address what an exporter must secure from an overseas recipient to satisfy the APPI. An organization contracting with an overseas AI provider engages both instruments at once, and 3 of 16 participants (19%) already reference METI's Contract Checklist for AI Utilization and Development. Making the two consistent, and indicating which governs where they meet, would let a single contract and a single internal review satisfy both.

OPPORTUNITY 5

For the PPC and METI, with industry:

> **Make provider adoption of the (forthcoming) model clauses visible, so exporters can verify rather than negotiate**

Model clauses relieve the constraints described by participants only if providers accept them. One participant reports that the terms an organization can obtain depend on its bargaining position and that the largest AI providers rarely negotiate. That participant is among one of the largest organizations in the cohort, which indicates the constraint reaches beyond smaller organizations.

Publishing the clauses in English alongside Japanese would encourage overseas providers to adopt them without first commissioning a translation. Maintaining a public record of providers that have adopted them unamended would also help to turn an exporter's task into checking a list rather than opening a negotiation in which it may have limited bargaining power.

Regional interoperability is what makes provider-side adoption commercially rational. A provider serving several Asian markets can adopt one clause set rather than negotiate market by market, and the exporters who gain most are those with the least individual leverage.

2.5 The rules that will govern wearable devices are being written now, and the decisive questions are open

Out of our 16 participants, only one develops or deploys advanced wearable devices for the Japanese market, and none manufactures wearable devices. Despite this, definitional clarity regarding biometric data ranked as participants' third highest policy priority in the survey. The APPI Amendment Bill introduced into the APPI a new category of "**specific biometric personal information**" (特定生体個人情報) that focuses on how the data is captured, rather than what that data is ultimately used for. This section considers two key questions that will influence whether this new category successfully supports the development of wearable technologies within Japan: whether the purpose for which data is processed affects the category's application; and whether a notice given for recording extends to applying recognition to that recording later. It also considers a constraint that sits outside the formal legal framework: whether the public accepts a particular device or sensing system being put into use in a particular setting.

2.5.1 The evidence base for this finding

This particular finding rests on a narrower evidence base than the others in our report.

15 of the 16 participants (94%) indicated in the survey that they are not currently developing, deploying, or planning to deploy advanced wearable devices for the Japanese market. This question functioned as a filter within our survey. It routed those 15 participants past the subsequent questions concerning the use of wearable data in AI, multimodal data collection, bystander privacy practices, and compliance challenges where the APPI intersects with other frameworks. Only one participant indicated that it currently deploys wearable technologies in Japan and consequently completed that section of the survey. Its responses are detailed in **§2.5.2** and are reflected throughout the data.

Therefore, to supplement the survey's findings, we solicited additional written input from four members of the program's Expert Group who possess relevant specialized knowledge. This group included one industry professional, one academic, and two legal practitioners based in Japan. Each expert received the same four core questions, alongside a fifth question specifically tailored to their individual field of expertise. Their responses are reported anonymously and are clearly identified as expert input wherever they appear in the text. These expert opinions are excluded from any quantitative figures presented in this chapter. We asked the experts for their independent assessments of the regulatory framework, rather than simply asking for comments on our preliminary findings. Whenever their perspectives diverge from what the participants told us, we explicitly highlight that difference.

2.5.2 What participants told us

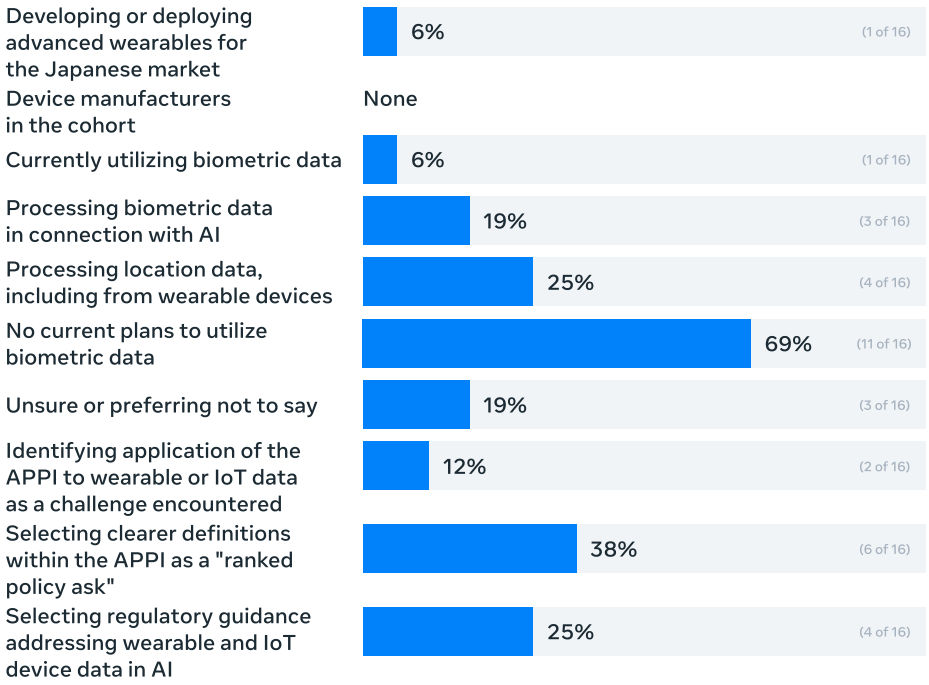


FIGURE 2.6
Engagement with biometric and body-based data (n=16)

Demand for guidance runs ahead of practice. Only 2 of the 16 participants (12%) have encountered a live problem when applying the APPI to data from wearable or Internet of Things (IoT) devices. Nevertheless, 6 of the 16 participants (38%) selected clearer definitions within the APPI concerning biometric data as a key policy change that would most benefit their operations. 4 of the 16 participants (25%) specifically selected regulatory guidance addressing wearable and IoT device data. 2 of the 16 participants (12%) ranked this guidance as their absolute highest priority among fifteen available options. One participant clarified that its ranking represented what it considered the Japanese framework needs generally, rather than reflecting a priority specific to the current needs of its organization.

The participant that does deploy wearable devices indicates that it processes names, personal identifiers, biometric data, and location data gathered from wearable devices. It handles personal information across three distinct stages of its AI workflow (which includes training and fine-tuning AI models). The multimodal data collected by its devices comprises gaze and eye-tracking data, along with video capturing visual surroundings and spatial data. It expects to begin using wearable device data for AI development within the next two years.

This participant's assessment of the framework is notably positive. In its responses to the survey, it somewhat agrees that the APPI offers clear guidance regarding the collection of personal information from non-user bystanders and that PPC guidance is sufficient for navigating privacy issues specific to wearables. Furthermore, the organization strongly agrees that it is operationally feasible to notify individuals in advance regarding the purpose of using “specific biometric personal information” collected via always-on environmental sensors.

The participant requested **specific guidelines covering wearable biometric data and AI-generated health inference data**. It asked that these guidelines distinguish clearly between data that is acquired, data that is inferred, and data that is used for training purposes, with separate consent requirements and usage conditions for each of those categories.

2.5.3 What the new category requires, and what remains open

The new category of “specific biometric personal information” is based on two cumulative criteria. First, it must be possible to capture an individual’s bodily feature **without special technology** (特別の技術) or **significant cost** (多額の費用). Second, the individual must be **unable to easily recognize** that this data capture is occurring. Further details will be specified in a forthcoming Cabinet Order.

The PPC has published a draft overview for implementation of the new category. In August 2026, the PPC's Secretariat issued a draft overview[6] outlining the specific regulatory matters that will be governed by forthcoming Cabinet Orders, PPC Rules, guidelines, and Q&A documents. Under a section concerning rules for facial feature data and similar information (顔特徴データ等に関する規律), the overview proposes three primary consequences for this new data category. First, organizations must make public (周知) certain information regarding

[6] https://www.ppc.go.jp/files/pdf/260717_kaiseihounitsuite.pdf

the handling of this data. Second, the legal conditions that an individual must meet to formally request the cessation of data usage have been relaxed, simplifying the process for individuals to bring such requests. Third, opt-in consent must be obtained before disclosing this data to third parties – reliance on the APPI's opt-out mechanism is prohibited.

The overview also contemplates establishing exceptions to the public knowledge obligation, and the relaxed right to request data cessation, as well as specific exemptions for handling that does not harm an individual's rights and interests because the collection is consistent with their underlying intention.

On present indications, the initial scope of the category appears to be limited to facial feature data. The experts observed that both the PPC's explanatory materials and the legislative discussions focused almost exclusively on collection of facial feature data using video cameras. Based on this, the experts assess that the regulatory scope will likely be limited strictly to facial features for the time being, leaving the inclusion of wearable devices uncertain. Consequently, the **experts recommend that the upcoming Cabinet Order explicitly state whether any other forms of specific biometric personal information fall within this new category.**

Experts identified the provisions on collection without special technology or significant cost as the principal source of difficulty. Regarding the "special technology" (特別の技術) provision, one expert considered that establishing a fixed definition of this term would be unrealistic as technology regarded as special today may become completely outdated within three to five years, new forms will continuously emerge, and regulatory authorities cannot update a definition continuously. That expert suggests adopting the approach taken in the European Data Protection Board's guidelines on data protection impact assessments for processing that is "likely to result in high risk."[7] The expert suggested that this framework's flexible concept of "new technology" could offer a workable model. A second expert, possessing a technical background, reached the same conclusion independently. That expert was fairly opposed to regulating the exact form of technology and instead, recommended that guidance should specify the functional aspects of the technology and the specific conditions under which it is used, coupled with a stated mechanism for updating the rules as new technologies appear.

The experts considered that the "significant cost" (多額の費用) provision presents even greater difficulties. One expert explained that costs are relative: an amount considered substantial for a small or medium-sized enterprise may not be substantial for a large company. Furthermore, prices fluctuate over time, meaning technology falling within the threshold today may fall outside it later – for instance, following a discount obtained on a contract renewal. In this expert's view, **the threshold cannot be expressed quantitatively**. Instead, it is critically important that regulators **specify the qualitative factors** relevant to determining where the boundary lies.

Practitioners facing challenges giving legal advice. One expert, referencing the PPC's August 2026 overview, observed that the specifics of every regulatory element remain entirely open to interpretation. Consequently, legal practitioners in Japan are encountering challenges in advising stakeholders under these circumstances.

[7] https://ec.europa.eu/newsroom/article29/items/611236

2.5.4 Three questions that will determine whether the framework supports development

① **Is a person who chooses to wear a device easily aware of the collection?**

One criterion for whether data qualifies as “specific biometric personal information” is whether an individual can easily recognize that capture is occurring. On its face, this could apply to two quite different populations – the person who has chosen to wear a device, and the people around them who have not – but the provision does not distinguish between them. This, in turn, determines the scope of obligations that apply to manufacturers of wearable devices and developers of applications that run on these devices and so shapes their design decisions taken years before a product launches.

Three experts consulted independently reached the same view on this question. They broadly agree that individuals who voluntarily wear a device and receive adequate disclosure about its data practices – such as through terms of use – could reasonably be considered easily aware of the collection.

The PPC's August 2026 overview identifies a further route to achieving this same practical outcome. The overview proposes an exception to the relaxed cessation right for cases where it is clear from the circumstances of acquisition that the data handling does not harm the individual's rights and interests. This applies specifically because the handling is not contrary to the individual's intention. A person who chooses to wear a device, having been explicitly told what it collects, is a natural candidate for this exact exception.

② **Does the purpose of processing matter?**

The new category of “specific biometric personal information” is scoped by how data is captured, rather than by what it is used for. While this appears to be a deliberate design choice, it distinguishes Japan's approach from frameworks in other jurisdictions that focus primarily on whether data is processed in order to identify a person.

The consequence of this scoping is that a great deal of processing in which no individual is identified may nonetheless fall within the category. Examples of

such processing may include gaze tracking for user selection in a display, gesture detection for broad user control of a digital interface, and in-cabin car functions that support alerting a driver during signs of fatigue.

The question matters now because the PPC's explanatory materials describe the category by reference to devices and software designed to identify individuals, while the statutory criteria do not mention purpose at all. Whether that framing carries through into the Cabinet Order and guidance is unresolved.

Here the expert responses diverge significantly, and the divergence provides valuable instruction.

- One expert considers that the regulatory concern appears focused heavily on devices and software specifically designed to identify individuals. This expert suggests it may be highly appropriate to clarify in the Cabinet Order or relevant guidance whether the active purpose of identifying an individual serves as a decisive regulatory factor.
- A second expert considers that temporary processing not intended to identify an individual will likely be considered separately from biometric data used specifically for identification purposes. However, this expert notes that such data will still constitute personal information, and therefore be subject to the APPI, if it can be linked to other identifying information about the individual.
- A third expert was less certain about purpose as a basis for categorization, on the ground that data which is retained can later be used for other purposes, and that it would be difficult for individuals or governments to confirm that a stated purpose continued to apply to current processing. On this view the more reliable control is architectural: developers should state the period for which data is retained and whether it leaves the device, which this expert expected would drive system design toward on-device processing.

③

Does a notice given for recording extend to applying recognition afterwards?

One participant raised this specific complex issue during their interview. On its account, obtaining baseline consent to be filmed and subsequently applying facial, gait, or other advanced recognition to that footage represent completely different layers of processing. The participant noted that these layers are not currently treated as distinct under the law.

An expert reached the same conclusion separately, arguing that it would be highly difficult to legally conclude that an individual's initial consent to standard video recording automatically extends to the subsequent application of facial recognition or other advanced recognition technologies. If the original explanation provided to the user specifically covered individual identification, the creation of facial feature data, and comparison against a database, then the later use may fall safely within its scope. However, if the individual was merely informed that they would be recorded, and the recording is subsequently used for individual identification or behavioral tracking, that processing may ordinarily exceed what the individual could reasonably have anticipated. This same expert notes that because specified information must be provided when collecting specific biometric personal information, a data subject will very likely need to be informed of those matters before giving valid consent. This required disclosure includes the fact of the processing, its exact purposes, and specific details of the bodily characteristics to be converted.

2.5.5 A non-regulatory constraint: public acceptance

When asked what most heavily influences whether organizations can successfully develop and deploy wearable devices in Japan, our experts identified a powerful factor completely outside the formal legal framework.

One expert described video recording of individuals as a topic of particular concern in Japan. The expert considered that this intense societal consideration must be built directly into the exceptionally high product quality standards that Japanese consumers expect. In this expert's assessment, organizations that consider legal compliance only after a product design is finished will inevitably produce products that fail commercially.

A second expert pointed to the exact same cultural history, noting that mandatory shutter sounds exist on Japanese smartphones directly because of public concerns regarding covert filming. Consequently, physical mechanisms such as an audible signal or a bright indicator light are universally expected of any devices that capture images in everyday settings.

A third expert perfectly summarized the unique social dynamic at play. In Japan, highly specific concerns initially raised by experts can spread incredibly quickly to the general public, rapidly developing into significant and overwhelming social criticism. This expert advised that conducting a rigorous risk assessment and executing a multi-stakeholder consultation process should be treated as practically mandatory, even though the APPI does not legally require these. The expert also considered that organizations should publish an accessible summary of these assessments, detailing the specific risks identified, the concrete mitigation measures incorporated directly into the product design, and highly practical guidance and frequently asked questions for everyday users.

This cultural context directly bears on how the remarkably limited industry engagement recorded in section 2.5.2 should be read. The regulatory position clearly remains unsettled, and participants are understandably waiting for it to settle. However, based on the experts' evidence, an organization that perfectly resolved every single regulatory question would still face a massive societal constraint fully capable of stopping a product deployment outright.

2.5.6 What this means for the ecosystem

This is not an unregulated space

The primary question facing Japanese policymakers is whether the existing regulatory interventions are calibrated to support technological development. The APPI currently applies to biometric data through the established individual identification code concept. The 2026 amendments formally add a dedicated category of specific biometric personal information. Sector-specific frameworks apply in health and medical contexts, while the AI Guidelines for Business offers guidance on the utilization of such data in AI systems.

The detail has yet to be provided, causing uncertainty

Three of the four expert responses converge on the exact same structural proposition, although they expressed it slightly differently. They assert that the Cabinet Order and PPC Rules should exclusively specify the operational conditions that must be satisfied and the qualitative factors relevant to establishing a regulatory boundary. Regulators must avoid prescribing specific forms of technology or creating fixed lists that will inevitably become outdated. One expert applied this logic directly to the issue of notice. This expert stated that the PPC should focus entirely on what functional conditions must be satisfied, rather than dictating how the requirements must be implemented. This approach intentionally creates space for private-sector actors to develop the most appropriate technological solutions.

The demand recorded in our survey should be read accordingly

Survey participants are explicitly asking the government to establish clear rules before they commit capital to build. Based heavily on the expert evidence, these organizations would benefit from a clear, principle-based statement explaining exactly what the new category covers, what specific operational criteria it turns on, and precisely how the boundaries will be assessed as underlying technology changes over time.

Recommendations for the PPC

RECOMMENDATION 4

Provide risk-based and workable definitional guidance on "specific biometric personal information" during the commencement period

We welcome the amendments to the APPI establishing a dedicated category for specific biometric personal information (特定生体個人情報). We also strongly welcome the PPC's publication of the overarching architecture by which the subordinate instruments to implement this new provision will be developed. Furthermore, we applaud its stated commitment to actively seek constructive views from individuals and businesses at successive stages of that drafting process.

Five specific clarifications would make this framework highly workable for organizations developing and deploying wearable technologies in Japan.

First, confirm that a person who chooses to wear a device, and is told what it collects, is easily aware of that collection

Every single legal expert we consulted independently reached this exact view. Future guidance would need to clearly address exactly what type of notice legally suffices, because awareness strictly follows from the individual fundamentally understanding what data is collected, rather than arising merely from the physical act of wearing the device. Accompanying guidance should also confirm that any data falling outside this new category firmly remains subject to the APPI's general obligations.

Second, clarify whether processing for the purpose of identifying an individual is a decisive factor

Based on the reading of our two legal experts, the PPC's primary regulatory concern appears to be directed at devices and software specifically designed to identify individuals. Formally confirming this interpretation would perfectly resolve the exact operational question that one of our participants could not answer when deploying wearable devices. It would also safely place completely benign processing where no individual is identified entirely outside the new category. This approach correctly leaves all such data strictly subject to the APPI's general obligations insofar as it can be linked to other identifying information about the individual.

RECOMMENDATION 4 CONTINUED



Third, state which modalities are within scope, and how boundary cases will be assessed as technology changes

Our experts assess that present regulatory indications suggest that implementation of the new category will initially focus on facial feature data, while the position of wearable devices remains unclear. Our experts recommend that the Cabinet Order state explicitly whether any other forms of specific biometric personal information fall within this category, so that industry does not rely solely on assumptions or inference. Additionally, purely quantitative criteria may be unsuitable for determining when capture requires special technology, and whether it would incur significant cost. As our experts explained, what counts as special technology changes quickly over time, and what qualifies as a significant cost depends on who is bearing that cost and when. Three of our four experts asked that the guidance specify the precise qualitative factors relevant to establishing these boundaries, alongside a formal mechanism detailing exactly how the assessment will be updated. One expert suggested that the flexible approach taken in the European Data Protection Board's guidelines on data protection impact assessments, which turn on the adaptable concept of "new technology," could serve as a workable model.



Fourth, address the relationship between capture and subsequent recognition, and the conditions under which the public knowledge (周知) obligation is engaged

Sensor data or video footage that is perfectly lawfully recorded may have advanced recognition algorithms applied to it later. Based heavily on the expert evidence, an initial notice of recording will not ordinarily extend legally to that later processing. Where the individual was merely informed that they would be recorded, later use for identification or behavioural tracking may exceed what the individual could reasonably have anticipated. It would therefore be useful if guidance stated what must be disclosed before collection, where later application of recognition technologies is contemplated – on one expert's account, this is likely to include the matters the Act already requires to be provided when specific biometric personal information is collected. The guidance must also definitively address whether any regulatory obligations are engaged at all in scenarios where data is strictly processed on the local device, deleted immediately after use, and never transmitted externally. One expert proposed this precise scenario directly as a valid condition under which all notice requirements could be safely relaxed. This approach represents the absolute clearest available route to making legal compliance a core property of a product's fundamental design, rather than reducing it to a mere paperwork disclosure exercise.

RECOMMENDATION 4 CONTINUED

Fifth, specify what operational conditions must be satisfied, rather than dictating precisely how they are to be technically met

Guidance of this kind works best when it states the outcome to be achieved and allows industry to determine how to achieve it. As one expert put it, this is not an area in which regulators should be expected to supply detailed technical guidance: the organizations building these systems understand them best, and are well placed to develop solutions suited to their own products. A second expert reached the same conclusion from a different direction. Technical requirements specific to Japan would add to device costs, which suggests that guidance framed around conditions rather than mechanisms would serve both regulatory and commercial objectives.

On process: These vital questions fall to be settled through the Cabinet Order, the PPC Rules, and the ongoing public consultation process. The primary risk identified by our comprehensive research is not under-regulation, but rather premature over-breadth. If rigid criteria are settled before commercial practice has properly matured, and drawn widely enough to mistakenly capture benign processing that presents absolutely no identification risk, the regulations would foreclose the exact innovative use cases the category was originally created to govern.

Opportunities for other policymakers

OPPORTUNITY 6

For the PPC and METI, in the commencement period:

> ## Establish the conditions under which incidental capture is adequately notified

The APPI requires certain matters to be made known when specific biometric personal information is collected. Where the person captured is the user of a device, that is straightforward: they can be told at the point of purchase or registration. Where the person captured is someone the device happens to record — a passer-by, or a person standing behind the user – the organization has no way to reach them, because it does not know who they are. Publication of a general notice to the public may satisfy this obligation formally; however, as one expert observed, an individual cannot know which published notice to consult when they do not know their data has been collected at all.

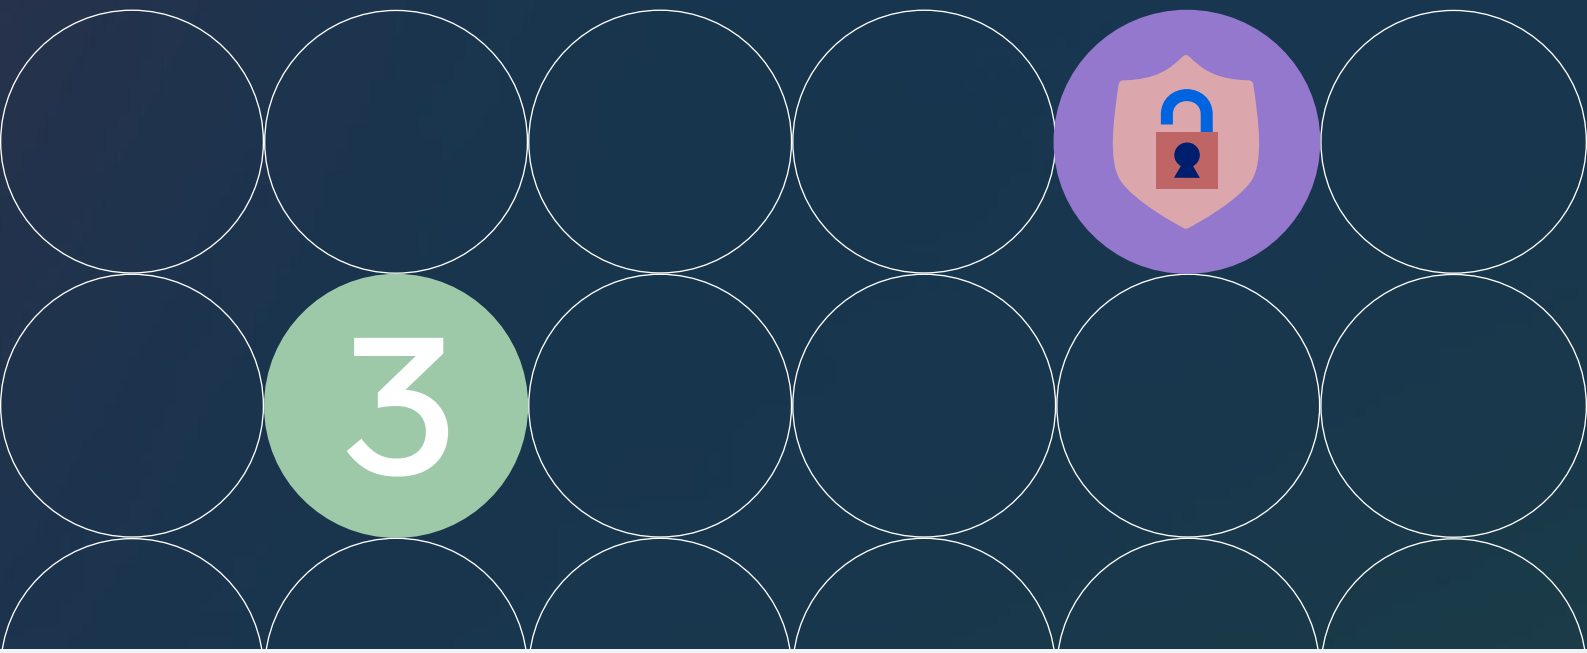
Rather than mandating specific technological mechanisms, the PPC and METI could jointly state the conditions that any proposed solution must successfully satisfy. These conditions should mandate that a person in the immediate vicinity can easily perceive that data collection is occurring, that they can readily obtain further detailed information if they wish, and that they possess a functional means of objecting. Private industry may be better placed to determine exactly how those broad conditions are technically met.

For instance, our experts described a wide range of innovative approaches already under active discussion, including explicitly displaying the device's operating status directly on the hardware itself. They suggested that developers might also use a visual indicator or a distinct audible notification recognizable by people nearby, repeated at regular intervals. Experts discussed a range of other potential future speculative approaches which could offer greater bystander transparency and control. Our experts' view was that regulators should state the conditions a solution must satisfy but then leave the technical design and implementation to industry.

One expert proposed a further condition highly worth stating explicitly. This expert suggested that where data is strictly processed on the local device, deleted immediately, and never transmitted externally, the mandatory notice requirements may not need to be engaged at all.

CHAPTER THREE

Conclusion & next steps



3.1 Summary of key takeaways

This program set out to understand how organizations operating in Japan are navigating the intersection of AI innovation and data protection, at a moment when the laws governing these areas were themselves changing. Sixteen organizations completed a structured bilingual survey; eight took part in follow-up interviews.

Between the release of the survey and the conclusion of interviews, the Diet passed significant amendments to the Act on the Protection of Personal Information (**APPI**). Two of those amendments bear directly on the questions addressed by this report.

- The first is a new exemption from consent for “statistical creation” (統計作成等), which permits both the acquisition of publicly disclosed sensitive personal information, and the provision of personal information to a third party for statistical purposes. Instead of requiring individual consent in these cases, the exemption requires organizations to publicize prescribed information about the statistical creation activity, on a continuing basis. According to the Cabinet's summary of the Amendment Bill, “statistical creation” extends to AI development.
- The second is a new category of “specific biometric personal information” (特定生体個人情報), determined by reference to whether a bodily feature can be captured without special technology or significant cost, and whether the individual can easily recognize that capture is occurring. This new category of personal information is subject to transparency requirements, stricter standards for disclosure to third parties, and a right for individuals to demand that use or provision cease, even where the information has not been handled unlawfully.

The amendments were promulgated on 17 July 2026, and most of the new obligations take effect by 16 July 2028. The practical scope of both provisions will be settled by Cabinet orders and PPC regulations that had not been published at the time of writing.



OBSERVATION 1

Industry expects the reform to help

12 of 16 participants (75%) anticipated that the 2026 amendments to the APPI would constructively reduce regulatory barriers to using personal information for AI – the largest measure of agreement recorded anywhere in our survey. Interviews conducted after the Amendment Bill passed found that expectation substantially unchanged. This report therefore arrives at a point in the reform process where industry feedback is most likely to be useful: the primary legislation is settled, the subordinate rules and guidance that will determine how it applies are still being made, and the organizations that will have to apply them are already working through what compliance will require.

OBSERVATION 2

Participants are not asking for lighter obligations

Across the survey and the interviews, participants primarily called for rules that they can apply before they act: criteria assessable in advance, examples, model contract clauses, and checklists. Where a participant asked for something more substantial, it was usually for the regulator to resolve questions of legal interpretation – such as whether a trained model constitutes a personal information database – rather than for approval to do something the rules currently prohibit. One participant asked for broader enforcement, not less.

OBSERVATION 3

Regulatory uncertainty is not stopping AI adoption, but it is shaping it

No participant rated regulatory uncertainty a "significant" or "primary" factor in their decisions concerning the development or deployment of AI. However, the report did find that regulatory uncertainty produces internal review costs, conservative responses, decisions taken without a settled view of their legality, and use cases that are narrowed or never attempted. The immediate consequence is that organizations forgo use cases for data that may have been lawful, because the organization cannot establish in advance that the law permits them. The broader consequence is that organizations reach different conclusions about the same requirement, with no mechanism by which those conclusions converge and no visibility for the regulator into where they differ.

Against that background, the five findings set out in Chapter Two are as follows.

FINDING 1

Japan's AI guidance reaches its intended audience, but participants do not find it sufficient to resolve questions of compliance with binding legal obligations

The AI Promotion Act places only best-effort duties on private organizations, so for an organization deciding whether an AI use case may proceed, the question carrying legal consequence is what the APPI requires. Japan's voluntary AI guidance is widely consulted – 11 of 16 participants (69%) take reference from the AI Guidelines for Business, and 3 of 16 (19%) treat them as their single most impactful framework. However, only one participant (6%) finds that guidance sufficient to establish what the APPI requires without further legal advice, and 14 of 16 (88%) either need legal advice on top of it or find that it does not resolve the questions they face. Each organization is left to reconcile the various governance frameworks on its own, and pays for that in internal guideline-writing, checklist-building and review time.

FINDING 2

Organizations have difficulty establishing how far they must de-identify personal information to use it in their AI activities

Practical, objective, risk-based standards for sufficient anonymization is the cohort's highest-ranked policy ask: 10 of 16 (62%) selected it, and 4 of 16 (25%) placed it first. 8 of 16 (50%) have encountered difficulties in practice. Two participants told us they have stopped relying on the statutory regimes altogether, treating all data as within the APPI's scope throughout – a rational response to a threshold that cannot be established in advance, but one that forgoes the flexibility that the APPI's anonymously processed information and pseudonymously processed information regimes were designed to provide.

FINDING 3

The classification of AI outputs and models under the APPI remains unresolved

8 of 16 participants (50%) cited the risk that generative AI outputs could reconstruct or infer personal information, placing it in joint first position among the compliance challenges participants had encountered. Beneath that sits an unresolved question about whether a trained model constitutes a personal information database under the APPI, which two participants raised independently. Where the question is unresolved, the more common response is restraint: organizations narrow the use case, or do not build it.

FINDING 4

Cross-border rules prevent participants from using AI systems hosted outside of Japan, and cloud rules cause uncertainty and hesitancy

7 of 16 participants (44%) report uncertainty about which overseas providers or jurisdictions satisfy the APPI's requirements, and this is the area in which participants describe having altered what they built. None of the 16 participants holds or is pursuing Global CBPR certification, yet 11 of 16 (69%) want confirmation that transfer to a certified recipient is lawful without any other basis. Since the PPC clarified that position in December 2025, this is a question of visibility rather than of law.

FINDING 5

The rules that will govern wearable devices are being written now, and the decisive questions are open

Only one of the 16 participants develops or deploys advanced wearable devices for the Japanese market, and none manufactures devices. Participants nonetheless rank definitional clarity about biometric data third among their policy asks, seeking rules before they build, rather than after encountering a problem. The 2026 amendments create a new category of specific biometric personal information, scoped by how data is captured rather than by what it is used for, and its content will be settled by Cabinet Order and PPC Rules over the commencement period.

The report identified three questions that will determine how the category applies to wearable devices, and additional input was sought from experts to answer them.

① Whether a person who chooses to wear a device is treated as easily aware of its data collection

Every legal expert we consulted considered that a wearer who has been told what the device collects is likely to fall outside the category.

② Whether the purpose of processing matters

Two experts read the PPC's explanatory materials as directed at devices designed to identify individuals, which would place routine data collection for non-identifying purposes (e.g., gaze tracking, gesture control, and fatigue detection) outside the category, while a third cautioned that purpose is an unstable basis for categorization because retained data can later be put to other uses.

③ Whether a notice given for recording extends to applying recognition technology to that footage later

An expert considered it would not, unless the original explanation specifically covered identification.

None of these three questions has been settled, and all three fall to implementing regulations and guidance for the 2026 amendments to the APPI that are currently being drafted. We recommend that the PPC resolve each during the commencement period, stating the conditions to be satisfied rather than the means of satisfying them, and setting out how boundary cases will be assessed as the technology changes (see **Recommendation 4**).

A fourth constraint sits outside the legal framework altogether. Our experts considered public acceptance of a particular device being put into use in a particular place more likely than any legal question to determine whether a product reaches the market.

The recommendations in **Chapter Two** are directed at the implementation of this reform. They identify the questions organizations will need answered before the new obligations take effect, and the form in which answers would be usable — published criteria, worked examples, model clauses, and Cabinet order provisions drawn narrowly enough to be applied at the design stage. They do not argue for different rules.

3.2 Implications for Japanese policy

These conclusions bear on Japanese policy in four respects:

What kind of instrument is needed	When it needs to arrive	What it needs to cover	Who needs to produce it

Almost nothing participants ask for requires further legislation

The questions identified in this report – how far data must be de-identified before it falls outside the APPI, what status attaches to what a model returns, how an exporter should assess an overseas provider, and how the new biometric category is scoped – are questions of interpretation. Through Cabinet orders and guidance, the answers can arrive within the period before the amended APPI takes effect, rather than awaiting a further review cycle.

The interval before commencement is the operative period

Participants describe waiting for subordinate rules and guidance before acting, which is a rational response to a reform whose application is not yet settled. However, a grace period spent waiting is not a grace period used. Guidance published early enough to be absorbed allows organizations to build to it; guidance published shortly before obligations take effect does not.

Several of the questions cut across institutional boundaries

Guidance organized around cameras and sensors rather than around industry sectors would require the PPC, METI, and MIC to act together. One participant observed that because responsibility for software, hardware and data policy sits with different ministries, guidance developed separately by each does not readily map onto products that combine all three.

A facilitative framework has to address development, not only deployment

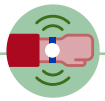
Much of this report concerns organizations applying rules to AI systems they are adopting. For those that develop connected devices, the earlier question is whether the framework supports building in Japan at all. This turns on matters arising well before a product reaches the market – for instance, whether a developer can determine at the design stage which obligations will attach to the sensor data generated by the device; whether the people whose data it captures will be treated as able to recognize that capture; whether development data collected from consenting participants can be used to train the models the device will run; and whether the answers will hold across a multi-year program. On the evidence of this cohort, none of those can currently be answered with confidence.

3.3 Future directions

What this research could not establish

Three questions remain open:

1. The first is whether difficulties described by participants are primarily a function of the rules or of organizations' capacity to interpret them.
2. The second is whether organizations reading the same requirement actually reach different substantive conclusions: we show that they organize the work differently and that some proceed on unresolved footings, but we did not put a common scenario to every participant, which would have tested the point directly.
3. The third is whether these requirements achieve their protective purpose – a question that cannot be answered by asking regulated organizations, and would require evidence from individuals, from complaint data, or from the regulator.



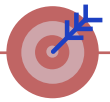
The evidence base on wearable devices

Our survey contained a filter question that routed 15 of 16 participants past the questions addressing wearable data, multimodal collection and bystander privacy. One participant passed the filter and answered them, and the interviews created opportunities to discuss these areas in further detail. Even so, Finding 5 rests on a narrower base of evidence than the report's other findings. Supplementary written input therefore had to be sought from members of the program's Expert Group with relevant expertise.



Regulatory milestones to monitor

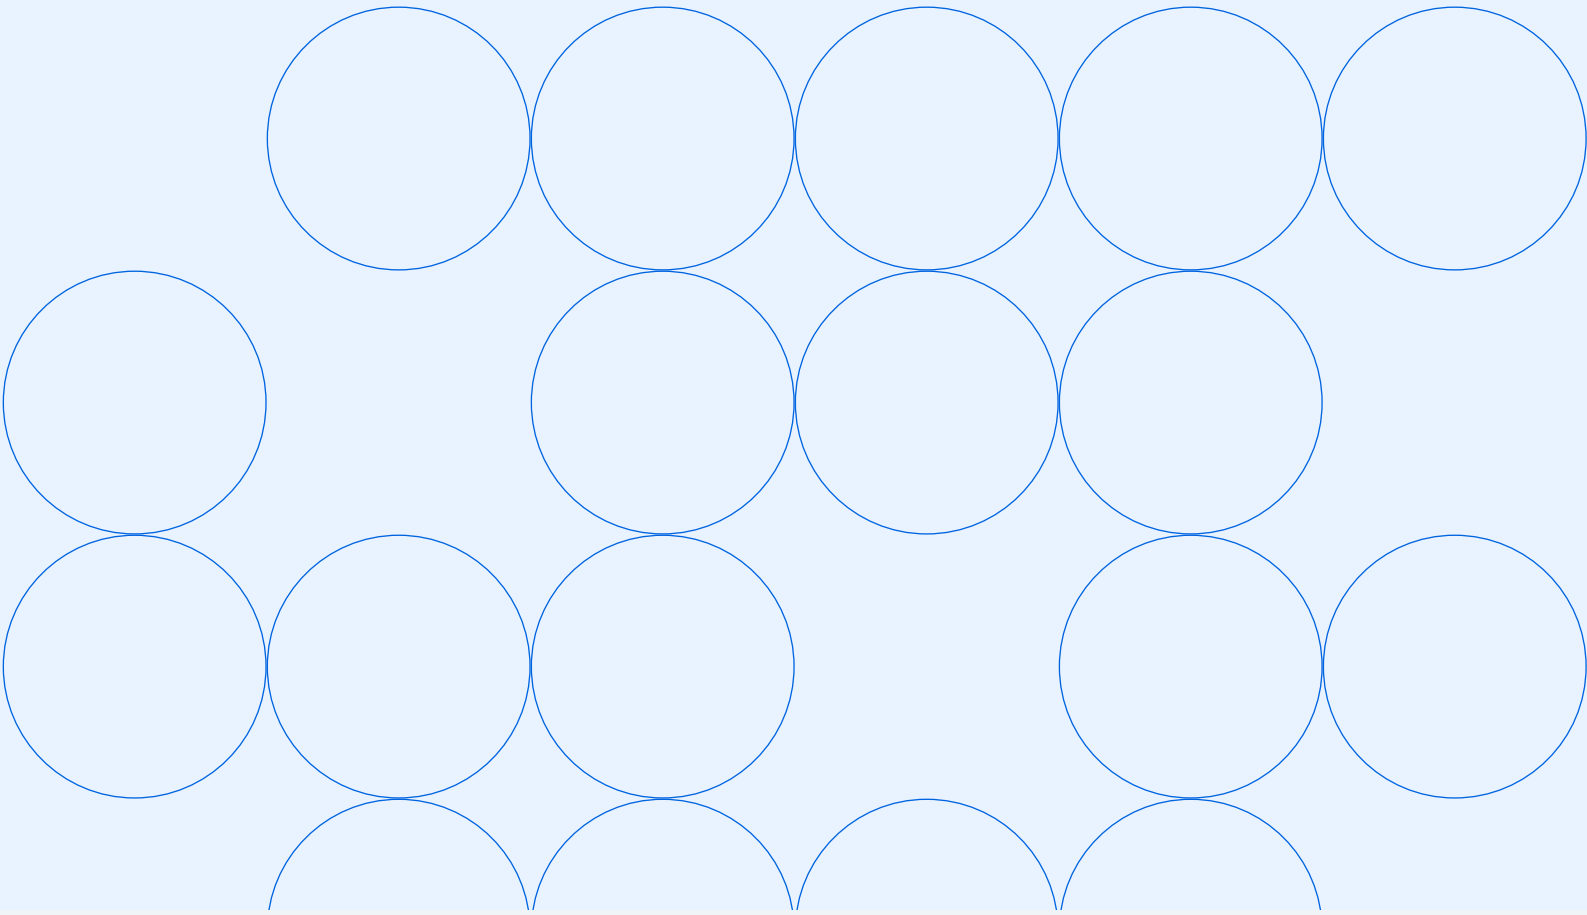
The Cabinet orders and PPC regulations that will scope both the new APPI provisions; the public consultations that accompany them; and the guidance issued ahead of commencement in July 2028. Each will determine how much of what participants describe is resolved. A follow-up exercise with the same cohort after commencement would establish whether the expectations recorded in this report were met – a question this report can pose but cannot answer at present.



Where further work would be most valuable

The status of trained models under the APPI's definitions, and what follows for the rights of disclosure, correction and suspension, is the question on which participants were least able to reach a settled position and on which the least guidance exists. The relationship between capture and subsequent recognition – whether a notice given in respect of recording extends to applying recognition technology to that recording – arises now, in ordinary commercial deployments, and will arise more sharply once the new category takes effect. And the position of a device's own user under the awareness criterion is the question on which the practical scope of that category most depends.

Annexes



I. About Open Loop Japan

i.i Program rationale

Open Loop Japan 2026 is a multi-stakeholder policy prototyping program designed to generate evidence-based insights that support the development of effective AI and data governance frameworks in Japan. The program operated from June to September 2026, convening companies, policymakers, academics, and civil society to collaboratively examine how emerging regulatory frameworks can be implemented in practice.

Japan is at a critical juncture in its approach to AI governance. The enactment of the AI Promotion Act in May 2025 established strategic direction for AI policy, while the ongoing triennial review of the Act on the Protection of Personal Information (APPI) presents an opportunity to modernize data protection rules in light of rapid technological change. As a leader in international AI governance, particularly through the G7 Hiroshima AI Process, Japan's approach to balancing innovation and protection has significant influence across the Asia-Pacific region and beyond.

Open Loop Japan brings the program's established "policy prototyping" methodology to this regulatory moment, bridging the gap between policy intent and practical implementation by testing how companies can operationalize new governance requirements in their AI systems and data practices.

i.ii Partners

The program was delivered through a partnership between:

Meta

PROGRAM CONVENOR

Meta's Open Loop initiative provided the program's global infrastructure, methodology, and convening platform.

FUTURE OF PRIVACY FORUM

LEAD RESEARCH PARTNER

FPF led the program's research design, data collection, expert group facilitation, and final report drafting.

accenture

LOCAL DELIVERY PARTNER

Accenture supported company recruitment, research planning, surveys, Tokyo event delivery, and translation.

i.iii Participating Companies (The Cohort)

HOW WE SELECTED PARTICIPANTS

Open Loop Japan recruited companies that develop, deploy, or use AI systems in Japan. Participants were identified through the networks of Meta, the Future of Privacy Forum, and Accenture, and through members of the project's Expert Group. The selection criteria were as follows:

- **Operations in Japan:** Each participant carries out AI-related activities in Japan that are subject to Japanese law – whether as a Japanese company or as the Japanese operation of an international group.
- **Direct experience of the questions under study**: Each participant applies the APPI, the AI Promotion Act, and/or Japan's voluntary AI guidance to real AI-related activities, rather than considering these frameworks in the abstract.

We recruited for variety across sectors, organizational size, and position in the AI value chain. We did not recruit for statistical representativeness, and the cohort should not be read as a sample of Japanese industry.

THE COHORT

Sixteen organizations completed a bilingual questionnaire of 25 questions on how they use AI, the personal information involved, the frameworks they consult, and the uncertainties they encounter. Eight of those organizations also took part in follow-up interviews. Further information on the survey and interviews is presented under **Methodology** below.

Each participant nominated a respondent with responsibility for AI governance, privacy, legal affairs, or digital transformation, who was able to answer for the organization as a whole.

ATTRIBUTION

Findings are reported in aggregate. This report does not attribute any statement, figure, or position to a named organization or individual and does not use sectoral descriptions from which an individual participant could be identified.

i.iv Methodology

RESEARCH APPROACH

This study uses a mixed-method design, combining a structured survey, semi-structured interviews, written input from subject-matter experts, and analysis of the relevant Japanese legal instruments and guidance.

DATA COLLECTION

Surveys: A survey was administered online in Japanese and English between 18 June and 25 August 2026. Each participant submitted one response, completed by the individual nominated to answer for the organization.

The survey comprised 25 questions covering the organization's role in the AI value chain, the personal information it handles and at what stage of its AI-related activities, its governance arrangements, the frameworks it consults, the uncertainties it encounters, the effect of those uncertainties on its decisions, and the policy changes it would prioritize. Questions combined single-select, multiple-select, ranked, and Likert formats with free-text fields.

- **Multiple-select questions**
  Participants could select more than one option on most questions, so the figures reported for those questions sum to more than 100%.
- **A filter question on wearable technology**
  Participants were asked whether they develop or deploy advanced wearable devices for the Japanese market. This question filtered participants past a subsequent block of wearable-specific questions. Because only a single participant answered the filter question in the affirmative, only that participant answered the subsequent questions. The evidence base for Finding 5 therefore rests principally on the interviews, input from subject matter experts (described further below), and the legal analysis, rather than on survey responses. **§3.3** discusses how this may impact future research.

Interviews: We conducted eight semi-structured interviews with participating organizations between August and September 2026. Interviews ran for approximately 60 minutes by video conference and were conducted in Japanese, with interpretation and transcription support.

Questions presented during the interviews were tailored to each participant's responses to the survey. The interviews

pursued three objectives: to establish the reasoning behind answers to the survey; to test the preliminary findings against the participant's own experience; and to ask what changes to policy in Japan (if any) may help to resolve questions or difficulties that the participant had raised in its responses to the survey.

Interviews conducted after the APPI Amendment Bill was promulgated on 17 July 2026 also allowed us to revisit participants' earlier expectations about the reform.

Expert Group: The research was supported by an Expert group of 20 members. The group comprised experts on technology governance and regulation from Japan and the broader Asia-Pacific region, drawn from academia, industry, the legal sector, and non-governmental organizations.

The Expert Group met virtually on three occasions over the course of the project. Not every expert attended every meeting. During the meetings, the Open Loop team shared information on the direction of the research and its preliminary findings, and experts asked questions and offered comments on the information presented. Several experts provided additional written feedback. As noted in the report, the Open Loop team also approached several experts for further written input on the questions raised in Finding 5 concerning wearable technologies and the APPI's new category of specific biometric personal information; four experts responded.

The views expressed in this report do not necessarily reflect the views of the experts.

ANALYSIS FRAMEWORK

Survey responses: We analyzed responses across the full cohort of 16 participants. Percentages are reported to the nearest whole number against a denominator of 16 throughout, including where a participant did not answer a given question. For the question asking participants to rank the policy changes that would most benefit their AI operations, we applied a weighted ranking that gives greater weight to higher placements, so that an option ranked first by a few participants is not displaced by an option ranked last by many. **§2.0** explains the distinction this chapter draws between those “ranked policy asks” and the separate question on framework priorities.

Interview material: We coded the interview transcripts against the themes emerging from the survey, and against the frameworks under study. Where an interview introduced a theme the survey had not anticipated, we returned to the survey data to test how widely it held. Where a participant's account conflicted with the survey data, we report the conflict rather than trying to resolve it.

Expert input: We treated the written expert responses as expert opinion on unsettled questions of Japanese law, not as consensus positions. Where experts disagreed, the report records the disagreement. Where an expert pointed to a primary source, we worked from that source: the discussion in **§2.5** of the criteria defining specific biometric personal information draws on the materials published by the PPC on 17 July and 26 August 2026.

Legal analysis: We analyzed the APPI as amended, the AI Promotion Act, the AI Guidelines for Business (Version 1.2), and the relevant PPC guidelines, Q&A, and published materials. Where the content of an obligation awaits a Cabinet Order, Commission Rules, or guidelines, we state this expressly.

Scope: Consistent with the project's focus, the analysis addresses the intersection of the APPI, the AI Promotion Act, and Japan's voluntary AI governance frameworks, and how that intersection bears on AI systems and wearable technologies. Matters raised by participants that fall outside this intersection are not treated as findings.

i.v Limitations

While this report offers valuable insights, its findings should be interpreted within certain methodological limitations:

Sample size and representativeness

The survey reflects responses from 16 organizations, which provides directional insights but limits the report's ability to generalize. The purposive sampling strategy was aimed at ensuring diversity across sectors and roles in the AI value chain; however, the results may not represent the full spectrum ofJapan's AI ecosystem.

Self-reported data

All findings are based on self-reported surveys and interview responses. These reflect participants' stated practices and perceptions rather than independently verified data.

Cross-sectional design

The study captures organizational practices and perceptions at a single point in time (mid-2026). As AI governance and compliance landscapes evolve, some findings may change.

No size-based disaggregation

The report does not differentiate responses by organization size (e.g., startups, small and medium enterprises, or large enterprises). This may obscure variations in compliance maturity or resource constraints that could influence responses. Several participants raised resourcing constraints directly, and **§2.0.3** and **§2.1.4** report those accounts, but the survey does not allow us to test how far the cohort's answers track organizational scale.

Potential social desirability bias

Considering the regulatory sensitivity of the topic, some responses may be influenced by reputational or compliance-related considerations. We note that several participants spoke candidly about unresolved questions affecting their own operations, which suggests this effect was not uniform across the cohort.

Limited cohort experience with advanced wearables

Only a single participant develops or deploys advanced wearable devices for the Japanese market. The survey's wearable-specific questions therefore produced little data, so the analysis in Finding 5 rests on the interviews, on written input from members of the expert group, and on the legal instruments themselves. **§3.3** sets out what this suggests for future research in this area.

Translation

The survey was administered in Japanese and English, and the interviews were conducted in Japanese, with interpretation and transcription support. Analysis was carried out in English. Where a participant's phrasing carried a technical or legal distinction, we have worked from the Japanese terms, which are given in the text where the distinction matters.

Subordinate instruments remain outstanding

Much of what the amended APPI will require of organizations working with AI and wearable technologies falls to Cabinet Orders, Commission Rules, and guidelines that had not been finalized at the time of writing. This report identifies the questions those instruments will settle and the considerations bearing on them, but it cannot state how they will be settled.

II. Key concepts and definitions

Agentic AI
Artificial intelligence (AI) systems that are capable of autonomous planning, decision-making, and action to achieve specific goals with limited human intervention.

Aggregation
The process of compiling and summarizing individual-level data into broader statistical formats (such as totals, averages, or ranges) so that information relating to specific individuals can no longer be isolated or identified.

Anonymization / Anonymously Processed Information (匿名加工情報)
Information relating to an individual that has been processed such that it is impossible to identify a specific individual or restore the original personal information. (Reference: Act on the Protection of Personal Information (APPI), Article 2).

Biomechanical data
Data relating to the mechanics of the human body (e.g., joint angles, muscle forces, and movement patterns), often collected via wearable devices or sensors as continuous streams of complex physical measurements.

Biometric data
Data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of an individual, which allow for that individual to be uniquely identified. Under the APPI, this falls under the category of an "Individual Identification Code" (個人識別符号), which includes data converted from physical characteristics for computer-based recognition (e.g., facial recognition data, fingerprints) (Reference: APPI, Article 2 and relevant Cabinet Orders).

Biometric inferences
The practice of using AI algorithms to analyze physical, physiological, or behavioral data (e.g., voice pitch, facial expressions, or eye movement), not to identify an individual but rather, to deduce secondary states or characteristics (e.g., emotional state, stress levels, or health conditions).

Bystander privacy
Privacy considerations regarding individuals who are not the primary users of a device, but whose personal information (e.g., face, voice, gait) may be incidentally captured by the device's outward-facing sensors (e.g., cameras, microphones) in public or shared spaces.

Pseudonymization / Pseudonymously Processed Information (仮名加工情報)
Information relating to an individual that has been processed such that it is not possible to identify a specific individual unless collated with other information.

(Reference: APPI, Article 2)

Sensitive Personal Information (要配慮個人情報)

Personal information as to an identifiable person's race, creed, social status, medical history, criminal record, the fact of having suffered damage by a crime, or other identifiers or their equivalent prescribed by Cabinet Order as those of requiring special care so as not to cause unjust discrimination, prejudice or other disadvantages to that person.

(Reference: APPI, Article 2)

Specific Biometric Personal Information (特定生体個人情報)

Personal information containing a specific biometric personal identification code, including those converted from information relating to physical characteristics of a part of the body that can be acquired by a method not requiring special technology or a large amount of cost, and for which the principal cannot easily recognize that the said information is being acquired.

(Reference: APPI Amendment Bill 2026, Article 16, Paragraph 5)

Wearable devices (Wearables)

Electronic devices designed to be worn on the user's body or incorporated into clothing (e.g., smart glasses, AR/VR headsets, smartwatches, and health trackers). These devices are frequently equipped with advanced sensors (such as cameras, microphones, and accelerometers) and edge computing capabilities, enabling the continuous collection and processing of both user-specific body-based data and ambient environmental data.

III. Endnotes

IV. References / bibliography