Generative Artificial Intelligence and Personal Data Protection Guidelines: Key Compliance Obligations for Companies Using Generative AI Systems

Authors: Dr. Gönenç Gürkaynak, Ceren Yıldız Özgökçen, Yasemin Doğan Yağcı and Asrin Özkahraman

I. Overview of the Generative Artificial Intelligence and Personal Data Protection Guidelines

On 24 November 2025, the Turkish Data Protection Authority (Kişisel Verileri Koruma Kurumu – "DPA") published the *Generative Artificial Intelligence and Personal Data Protection Guidelines* (the "**Guidelines**"). The document represents the DPA's first comprehensive attempt to address the data protection implications of generative artificial intelligence ("**Generative AI**") systems under Law No. 6698 on the Protection of Personal Data ("**DP Law**").

Rather than introducing new binding rules, the Guidelines aim to interpret existing personal data protection principles in light of rapidly evolving Generative AI technologies. In doing so, the DPA examines how personal data is processed at each stage of a generative AI system, from the initial development and training phase to deployment and day-to-day use. The Guidelines place particular emphasis on transparency, accountability, data minimization, and the protection of data subjects' rights. This reflects concerns that Generative AI models rely on large-scale data processing but often operate as "black boxes" offering little transparency about their decision-making processes.

Importantly, the Guidelines make clear that the novelty or technical complexity of Generative AI does not exempt companies from their obligations under the DP Law. On the contrary, the DPA signals that heightened risks associated with Generative AI systems require a correspondingly higher level of diligence from data controllers.

II. Scope of Application: Companies and Sectors Affected by the Guidelines

The scope of the Guidelines is intentionally broad. They apply to any natural or legal person that determines the purposes and means of processing personal data within the context of Generative AI systems. This includes not only developers of Generative AI models, but also companies that deploy, integrate, or use such systems in their business operations.

From a sectoral perspective, the Guidelines are relevant across a wide range of industries. While technology companies and AI developers are obvious stakeholders, the DPA expressly acknowledges the increasing use of Generative AI tools in areas such as finance, healthcare, education, media, marketing, and customer services. Even organizations that rely on third-party

Generative AI tools, rather than developing their own models, may qualify as data controllers if they determine how and for what purposes personal data is processed through those tools.

The Guidelines also highlight the importance of correctly identifying the roles of data controller and data processor within complex AI value chains. Depending on the specific phase of the AI lifecycle, different actors may assume distinct roles, and these roles may shift over time. Companies are therefore expected to conduct a careful, case-by-case assessment rather than relying on generic contractual labels.

III. Lawful Bases for Processing Personal Data in Generative AI Systems

Training and Deployment of Generative AI Models

One of the most sensitive issues addressed in the Guidelines concerns the lawful bases for processing personal data during both the training and deployment stages of Generative AI models. The DPA underlines that personal data may be processed at multiple points in the lifecycle of a Generative AI system, including the collection of training datasets, user interactions, and output generation.

The Guidelines emphasize that each processing activity must independently satisfy one of the lawful bases set out in Article 5 of the DP Law. In practice, this creates significant challenges, particularly where models are trained on large, heterogeneous datasets sourced from publicly available content. The DPA makes it clear that the mere public availability of data does not automatically render its use lawful.

Consent is addressed with notable caution. Given the scale, complexity, and future unpredictability of Generative AI systems, the DPA questions whether consent can genuinely be considered “informed” and “freely given” in many of the training scenarios. As a result, companies are encouraged to assess alternative lawful bases, such as the necessity of processing for the establishment or performance of a contract, or the existence of a legitimate interest—provided that such interest does not override the fundamental rights and freedoms of data subjects.

During deployment, lawful basis assessments must also cover the processing of user inputs, prompts, and potentially personal data contained in generated outputs. The Guidelines stress that companies cannot rely on a single, blanket legal basis to cover all stages of Generative AI usage.

IV. Key Compliance Obligations for Data Controllers Using Generative AI

Transparency and Information Obligations

Transparency is a recurring theme throughout the Guidelines. The DPA considers transparency to be particularly critical in the context of Generative AI, where data subjects may be unaware that their personal data is being processed or that AI-generated content is influencing decisions that affect them.

Data controllers are expected to provide clear, accessible, and meaningful information about how personal data is processed within Generative AI systems. This includes information about

the types of data processed, the purposes of processing, the logic of the system to the extent possible, and potential risks for data subjects. The DPA recognizes that full algorithmic explainability may not always be feasible, but it expects controllers to make genuine efforts to avoid vague or generic disclosures.

Data Minimization and Purpose Limitation

The Guidelines reaffirm the core principles of data minimization and purpose limitation, while acknowledging the tension between these principles and data-hungry Generative AI models. The DPA explicitly rejects the notion that the technical needs of AI systems justify indiscriminate data collection.

Data controllers are expected to critically assess whether personal data is truly necessary at each stage of the AI lifecycle. Where possible, anonymized or synthetic data should be preferred, particularly during training and testing phases. The DPA also warns against repurposing personal data for AI training, where such use would be incompatible with the original purpose of collection.

Technical and Organizational Measures

Given the elevated risks associated with Generative AI, the Guidelines place strong emphasis on technical and organizational safeguards. These include measures to prevent unauthorized access, data leakage, and the unintended memorization or reproduction of personal data in AI outputs.

The DPA encourages the adoption of privacy-by-design and privacy-by-default approaches, ensuring that data protection considerations are embedded into system architecture from the outset. Regular risk assessments, internal governance frameworks, and staff training are also highlighted as essential components of compliance.

Risks Relating to Invalid Consent and Unlawful Data Processing

The Guidelines repeatedly caution against over-reliance on consent as a lawful basis for Generative AI-related processing. In the DPA's view, consent obtained through lengthy, complex, or ambiguous information notices is unlikely to meet the validity requirements under the DP Law.

Invalid consent exposes companies to significant legal risks, including the risk that all subsequent processing activities may be deemed unlawful. This is particularly problematic in Generative AI contexts, where personal data may be deeply embedded within model parameters and difficult to extract or delete retroactively.

The DPA also draws attention to the risk of "function creep," whereby data initially collected for limited purposes gradually becomes embedded in AI systems used for broader or unrelated objectives.

Potential Regulatory and Enforcement Risks under the DP Law

While the Guidelines themselves are not legally binding, they provide a clear indication of the DPA's enforcement priorities. Failure to comply with the principles outlined in the Guidelines may increase the likelihood of administrative investigations, corrective orders, and administrative fines under the DP Law.

The DPA also signals that Generative AI systems may give rise to cross-border data transfer issues, particularly where cloud-based infrastructure or foreign AI service providers are involved. Non-compliance with international transfer rules may therefore compound regulatory exposure.

Practical Compliance Considerations for Companies Implementing Generative AI Tools

For companies seeking to implement or use Generative AI tools, the Guidelines point toward several practical steps. These include mapping AI-related data flows, identifying controller-processor roles with precision, conducting impact assessments where appropriate, and reviewing vendor contracts to ensure adequate data protection safeguards. Equally important is the establishment of internal governance mechanisms capable of responding to evolving risks.

Generative AI is not a static technology, and compliance efforts must be revisited as systems evolve, new use cases emerge, and regulatory expectations continue to develop. Accordingly, a proactive and adaptive approach to compliance will be essential for companies to navigate the shifting landscape of Generative AI regulation with confidence and accountability.

Article Contact: Dr. Gönenç Gürkaynak E-mail: gonenc.gurkaynak@elig.com

(First published by Mondaq on February 10, 2026)