

Public Comment on the EDPB DPIA Template (v1.0, March 2026): Three Gaps That Matter for Agentic AI Systems

Submitted to the EDPB DPIA Template public consultation (deadline 9 June 2026)

Author: Amine Raji, PhD, CISSP — AI security practitioner, Gothenburg, Sweden

Also published at <https://aminrj.com/edpb-dpia-template-comment>

The harmonised DPIA template adopted by the European Data Protection Board on 10 March 2026 is a substantive step toward consistent GDPR application across the EEA. After eight years of fragmented national approaches, a common minimum documentation standard is genuinely useful. The methodological distinction between design-structural risks (Section 3.1 of the Explainer) and deviation or incident risks (Section 4.1.1 of the Explainer), unified in the inherent risk assessment at Section 4.1.3, is a welcome refinement that will outlast version 1.0.

This comment focuses narrowly on one question: how well does the template, as drafted, accommodate Data Protection Impact Assessments for agentic AI systems? My background is in the technical security of agent architectures — Model Context Protocol servers, Retrieval-Augmented Generation pipelines, multi-agent orchestration — and I have run impact analyses against several deployments of these systems, both in lab environments and as part of enterprise security reviews.

The short answer: the template is workable for traditional data processing operations but leaves three gaps that, if not addressed, will make EU-wide DPIAs for agentic AI systems either incomplete or inconsistently structured across supervisory authorities. The Spanish AEPD's February 2026 guidance on agentic AI from the perspective of data protection — among the most detailed regulatory treatments of these systems by any EU DPA to date — identifies many of the same gaps. The Dutch Data Protection Authority issued a similar warning in February 2026 about highly autonomous AI agents. There is now a regulatory pattern that the harmonised DPIA template should reflect.

Note on section references: all section numbers in this comment follow the numbering used in the Explainer document. The template form (.docx) uses letter suffixes for sub-sections (e.g. 1.1.a, 2.3.b) where the Explainer uses numeric ones (e.g. 1.1.1, 2.3.2); the substantive fields are identical.

I respectfully submit three concrete proposals.

Gap 1: Section 1.3 ("Means of processing, supporting assets and underlying architecture") treats agent memory and tool-call interfaces as ordinary assets

Section 1.3 of the template asks the controller to list "essential supporting assets" and gives examples including "hardware, infrastructure, and network assets," "software (database engines, business applications)," "APIs and models," and "personnel." This taxonomy works well for

traditional processing pipelines. For agentic AI systems, it underspecifies two asset categories that are central to how the processing actually occurs.

Agent memory. Agentic systems maintain short-term and long-term memory that stores context across sessions, user preferences, and prior interactions. The AEPD's guidance specifically flags persistent memory as a significant compliance risk and requires that memory be compartmentalised between processing activities and users, with strict retention periods. None of this is captured by listing "database engines" as an asset. Memory in agentic systems is structurally different: it is read into the model context window at each invocation, it shapes downstream outputs in ways the controller cannot fully predict *ex ante*, and it can contain personal data that the data subject never provided directly (for example, inferences the agent made and stored).

Tool-call interfaces. Modern agent architectures expose toolkits — calendars, email systems, databases, internal services, third-party APIs — that the agent invokes autonomously to fulfil its goals. Each tool call is a data flow. A single user prompt can trigger a chain of dozens of tool calls, each potentially involving personal data, with no human decision at any individual step. The AEPD's guidance converges on a point that multiple legal commentators have also identified: these tool-call chains amplify vulnerabilities in ways that are not captured by traditional data-flow diagrams.

Proposed addition. Add a sub-section 1.3a or supplementary guidance in the Explainer for "AI-specific assets," with two non-exhaustive sub-categories: (a) "Memory systems," requiring the controller to identify each memory store, its scope (per-user, per-session, persistent), its retention behaviour, and which personal data categories it may contain; and (b) "Tool and service integrations," requiring the controller to enumerate the tools the agent can invoke, the personal data accessible through each, and the conditions under which the agent may call them. This change adds approximately one page of structured fields and gives DPAs a consistent basis for review.

Gap 2: Section 1.2 ("Functional description") assumes a linear data lifecycle that does not exist in agentic systems

Section 1.2 of the template asks the controller to "explain the data lifecycle and data flows" and structures this around five stages: Collection, Use, Storage, Sharing and Transfer, Deletion and Destruction. This is the canonical lifecycle that has anchored DPIAs since the Article 29 Working Party's 2017 guidelines (WP248), and it works well for processing operations where data moves through discrete stages.

It does not describe agentic processing accurately. In an agent architecture, data is not collected once and then used: data is retrieved continuously into the model context window, mixed with system prompts, recombined with retrieved documents, and passed into tool calls — frequently within a single user session. The AEPD calls this the "chain of reasoning" and warns that insufficient governance of the multiple processing purposes embedded in this chain can violate Article 5(1)(b) purpose limitation. RAG pipelines specifically retrieve personal data from indexed

corpora at inference time, which means the same personal data is processed for materially different purposes depending on the user prompt — a fact the linear lifecycle framing obscures.

A controller filling out Section 1.2 today, in good faith, can write a description that satisfies the form fields but does not reflect what the agent actually does. The supervisory authority reviewing the DPIA will not be in a position to identify the gap unless the template structure surfaces it.

Proposed addition. Either expand Section 1.2's prompts to explicitly require description of (a) inference-time retrieval (RAG and similar mechanisms), (b) context-window assembly (what data is concatenated at each invocation and from which sources), and (c) the chain of reasoning (how a single user prompt decomposes into sub-tasks and tool calls); or add an Annex specifically for AI-augmented processing operations that supplements the linear lifecycle with these additional descriptions. The AEPD's analytical framework — examining uncontrolled information processing, sensitive information access, and external action as distinct risk dimensions — provides one concrete structuring option.

Gap 3: The template is silent on the AI Act Article 26(9) interface

Article 26(9) of the AI Act, applicable from August 2026, requires deployers of high-risk AI systems to use the information provided by the AI system provider under Article 13 of the AI Act to fulfil their DPIA obligation under Article 35 GDPR. Article 27 of the AI Act creates a parallel Fundamental Rights Impact Assessment (FRIA) obligation for the same systems, and Article 27(4) explicitly allows deployers to leverage their existing DPIA for the FRIA, with both potentially consolidated into a single integrated report.

This means that from August 2026, a meaningful proportion of high-risk DPIAs will be conducted in parallel with FRIAs and will draw on AI Act Article 13 documentation. The current EDPB template makes no reference to this interface. Section 0.5 lists "Guidelines, standards, codes of conduct, and other reference materials" but does not prompt the controller to indicate whether AI Act Article 13 documentation has been received and used, nor whether the DPIA is being conducted in coordination with a FRIA.

This is not a minor cross-reference. The legal commentary on the FRIA-DPIA interplay — A&O Shearman's December 2025 *Zooming in on AI* series, the ECNL/DIHR FRIA guide, the Mantelero academic framework — converges on the view that consolidated DPIA-FRIA documents will be common practice. The EDPB's template, which is positioned to become the meta-template all national DPAs align with, sets the structural expectation that will either accommodate this practice or push controllers to maintain parallel uncoordinated documents.

Proposed addition. In Section 0.5 (DPIA technical sheet), add explicit checkbox prompts: (a) "Is this DPIA conducted in connection with a high-risk AI system under EU AI Act Annex III?" (b) "Has AI Act Article 13 documentation from the AI system provider been received and incorporated?" (c) "Is this DPIA part of an integrated DPIA-FRIA document under AI Act Article 27(4)?" The Explainer should reference the EDPB-AI Office interplay guidelines currently in

development, and should signal that the template may be updated once those guidelines are finalised.

A note on data subject rights at Sections 2.3.2 and 5.2

A brief additional observation. Section 2.3.2 lists measures to support data subject rights and Section 5.2 covers the views of data subjects or their representatives. For agentic systems specifically, exercising rights under Articles 15, 16, 17, and 22 is genuinely difficult: a request to access or erase personal data must reach not only the primary database but the agent's memory stores, any logs of tool calls that captured the data in transit, and any inferences the agent generated and persisted. The AEPD highlights this as a defining compliance risk for agentic deployments. The template would benefit from at least a note in the Explainer pointing controllers to the specific complications memory and chain-of-reasoning architectures introduce for rights exercise.

Closing

The harmonised template is a substantial improvement over the fragmented status quo, and the methodological clarity of the design-risk versus incident-risk distinction is a contribution that will outlast version 1.0. The proposals above do not require restructuring the template. They require modest additions — one supplementary asset sub-section, expanded prompts in the functional description, three checkboxes in the technical sheet, and a few paragraphs of Explainer guidance — that bring agentic AI within reach of the existing framework rather than leaving controllers to improvise.

If the EDPB finds these proposals useful and would benefit from concrete drafting, I would be glad to contribute. The Spanish AEPD has demonstrated that detailed regulatory engagement with agentic systems is feasible; the EDPB template should reflect what its own member authority has already published.

Respectfully submitted,

Amine Raji, PhD, CISSP

AI security practitioner — Molnsys AB, Gothenburg, Sweden

<https://molntek.com>

Contributor, OWASP Top 10 for Agentic Applications

amine [at] aminrj.com — <https://aminrj.com>

References

1. European Data Protection Board, "Template [2026] for Data Protection Impact Assessment ('DPIA') v1.0 (10 March 2026) and accompanying Explainer document.

2. Agencia Española de Protección de Datos (AEPD), "Inteligencia artificial agéntica desde la perspectiva de la protección de datos" (18 February 2026), 81 pp.
3. Article 29 Working Party, Guidelines on Data Protection Impact Assessment (wp248rev.01).
4. Regulation (EU) 2024/1689 (the EU AI Act), Articles 13, 26(9), 27, and 27(4).
5. A&O Shearman, "Zooming in on AI #13: EU AI Act – Focus on fundamental rights impact assessment for high-risk AI systems" (December 2025).
6. European Center for Not-for-Profit Law and Danish Institute for Human Rights, "A Guide to Fundamental Rights Impact Assessments (FRIA)" (December 2025).
7. A. Mantelero, "The Fundamental Rights Impact Assessment (FRIA) in the AI Act: Roots, legal obligations and key elements for a model template," *Computer Law & Security Review* (2024).
8. Dutch Data Protection Authority, public statement on autonomous AI agents (February 2026).