

The Patient Data Supply Chain Behind Your AI Service Contract

A Healthcare Board Decision on Actual Data Use

Phil Hatch | Akholi

October 23, 2026

Healthcare boards should require management to map patient-data use across the AI service supply chain. The contract must reflect actual recipients, purposes, memory, tools, support access, retention, and further use rather than describe only the prime supplier.

Approve the Processing That Actually Occurs

The board should require a clear account of where patient information goes and why before approving a material outsourced AI service. A prime supplier can rely on model hosts, tool providers, cloud platforms, support teams, and further subcontractors. The relevant question is the actual processing relationship, not the number of signatures on the main agreement.

HHS guidance explains business-associate and subcontractor duties where HIPAA applies.¹ It is a US framework with defined scope and exceptions, not a universal patient-data contract. The buyer should determine the applicable roles and obligations for the actual service and entities.

Begin with one patient-related workflow. Identify information collected, derived, retrieved, stored, disclosed, or retained in agent memory. Include debugging, human support, incident investigation, and product improvement. A statement that the model does not train on customer data does not explain these other uses.

¹ US Department of Health and Human Services Office for Civil Rights, “Business Associates Frequently Asked Questions (FAQs),” US Department of Health and Human Services Office for Civil Rights, accessed October 23, 2026, “What is a Business Associate” and “Business Associate Agreements,” <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/index.html>.

Separate Each Purpose and Recipient

The provider should identify which parties act only on instructions, which determine further purposes, and which receive information under another legal relationship. Technical integration does not settle those roles. EDPB guidance bases controller and processor analysis on the actual processing relationship.²

The study reports incomplete downstream dependency mapping in banking engagements.³ It does not establish healthcare privacy breaches. The transferable concern is whether management knows the recipients and operating dependencies that sit below the prime supplier and can obtain evidence about them.

For a hypothetical scheduling service, patient context may appear in prompts, logs, memory, analytics, and support tickets. The board should expect management to distinguish the purpose and retention of each copy. An approved scheduling purpose should not silently become permission for unrelated shared-model improvement.

Use a Patient Data Supply Chain Record

Figure 1 proposes a record linking purpose, recipient, operating control, and lifecycle treatment. Identify the data categories, locations, access roles, relevant agreement, retention, deletion, and incident route. Record uncertainty where the provider cannot establish a material fact.

The record should cover actual agent tools and memory rather than rely solely on a static architecture diagram. Test representative execution and support access. A new tool or changed memory design may introduce a recipient or purpose that was not assessed when the main contract was signed.

Management should secure appropriate evidence, authorization, and cooperation from material downstream parties. A general flow-down clause should be supported by a demonstrated route for investigating access, restricting use, correcting information, and handling required notifications. The board needs to know where that route remains incomplete.

² European Data Protection Board, *Guidelines 07/2020 on the Concepts of Controller and Processor in the GDPR* (European Data Protection Board, adopted July 7, 2021; version 2.1 corrected September 20, 2022), version 2.1, executive summary and part I, secs. 2–3, https://www.edpb.europa.eu/system/files/documents/2023-10/EDPB_guidelines_202007_controllerprocessor_final_en.pdf.

³ Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 3, Third-Party Ecosystem, sec. 6.1, “Mapping of Third and Fourth Party Dependencies,” indicator C196, <https://doi.org/10.2139/ssrn.7164559>.

Figure 1. The patient data supply chain record

Processing layer	Board question	Management evidence
Purpose	Why is each copy used	Service support memory and reuse map
Recipient	Who controls or receives information	Roles locations and downstream agreements
Control	Can use be investigated and restricted	Access evidence and incident cooperation
Lifecycle	What changes or ends processing	Retention deletion transfer and change tests

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Apply the Relevant Privacy and Health Duties

Swiss privacy guidance confirms that data-protection law applies to AI-supported processing.⁴ EU, UK, Swiss, US, and other healthcare services need separate assessment of lawful grounds, roles, secrecy, sensitive information, individual rights, transfers, and sector requirements. The same contract label should not be assumed to solve every jurisdiction.

Ownership of data is not the same as permission to process it. A commercial right to retain a workflow artifact does not automatically permit retention of patient information within it. Privacy and legal specialists should assess proposed reuse, de-identification claims, and retention against the actual data and purpose.

The incident method should identify who detects, preserves evidence, contains further use, and supports the responsible entity's legal decisions. Contract notice should supply the facts needed for those decisions without being mistaken for the applicable statutory deadline or notification itself.

⁴ Federal Data Protection and Information Commissioner, "AI and Data Protection," Federal Data Protection and Information Commissioner, September 24, 2025, "Transparent information" and "Protection of fundamental rights," <https://www.edoeb.admin.ch/en/ai-and-data-protection>.

Require a Current Map Before Expanding the Service

A reasonable objection is that the supply chain changes too quickly for a static review. That is a reason to define change triggers and maintain the record, not to omit the mapping. Prioritize material recipients and uses, and require timely notice and reassessment when the processing changes.

The board should receive the processing map, role decisions, material downstream limitations, tested rights and incident routes, and unresolved risks. Restrict new purposes or recipients where the required basis and controls are not established. A provider assurance should be evaluated within its actual scope.

Approve the AI service when the patient-data relationship is understood and governed across its full chain. The prime contract is the starting point. The operating evidence must show that the purposes, recipients, access, and lifecycle treatment it describes remain true as the service runs and changes.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. Its application here is a mechanism to investigate, not an estimate of conditions in other industries. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁵

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is a patient-data supply-chain record connecting actual agent execution and memory to processing purposes, downstream roles, operating evidence, and lifecycle decisions. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

European Data Protection Board. *Guidelines 07/2020 on the Concepts of Controller and Processor in the GDPR*. European Data Protection Board, adopted July 7, 2021; version 2.1 corrected September 20, 2022. https://www.edpb.europa.eu/system/files/documents/2023-10/EDPB_guidelines_202007_controllerprocessor_final_en.pdf.

Federal Data Protection and Information Commissioner. “AI and Data Protection.” Federal Data Protection and Information Commissioner. September 24, 2025. <https://www.edoeb.admin.ch/en/ai-and-data-protection>.

Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.

US Department of Health and Human Services Office for Civil Rights. “Business Associates Frequently Asked Questions (FAQs).” US Department of Health and Human Services Office for Civil Rights. Accessed October 23, 2026. <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/index.html>.

⁵ Hatch, *Agentic AI Outsourcing*, Introduction, “Executive Summary” and “Evidentiary Basis and Its Limits.”