

The State of AI Change Obligations in Healthcare

August 7, 2026 · Dilip Adityan, Founder · Gatekeeper (K2 Corp d/b/a Gatekeeper) · v1.0

Every load-bearing claim in this report carries a source. Where our own research corrected an earlier number, the correction is stated in the text — that discipline is the point of the report.

THE ONE-PARAGRAPH VERSION

A health-tech vendor's AI now changes on its platform providers' schedule, not its own — roughly one obligation-triggering event every three weeks across OpenAI, Anthropic, Azure, and Google. The health sector's own institutions have already written down what vendors owe their customers when that happens: 30 days' advance notice, release notes, revalidation evidence, incident clocks measured in hours. Yet no standalone AI addendum instrument exists anywhere in SEC EDGAR's full-text corpus going back to 2001, and the leading neutral contract templates contain no model-change notice duty at all. The demand side has published its spec. The instrument to satisfy it does not publicly exist. This report documents both halves of that gap — and publishes an open, machine-readable change-notice schema as a first attempt to close it.

1. The recurrence engine: your AI changes whether you like it or not

We assembled a trailing-12-month feed (August 2025 – August 2026) of exogenous, obligation-triggering platform changes from the four major AI providers' official lifecycle pages — model deprecations, retirements, breaking API changes, and forced migrations.

The number: 19 verified events — one every ~19 days. Collapse simultaneous same-day announcements the way a vendor's change-control queue would actually experience them and the count is 17, one every ~21 days. Either way: **roughly one obligation-triggering platform event every three weeks, imposed by the labs regardless of the vendor's roadmap.**

Three honest caveats, because they are findings in themselves:

- **Azure OpenAI contributed zero countable events** — not because nothing changed, but because Microsoft publishes retirement dates without announcement dates, making notice periods uncomputable from public sources. Microsoft's own two sources also give conflicting retirement dates for the same model (o3-mini: 2026-08-02 in the docs repo, 2026-10-01 on the rendered page). A vendor tracking that model got two deadlines from one publisher.
- **Subprocessor changes are uncountable.** Neither OpenAI nor Anthropic publishes a dated public subprocessor changelog — the only way to know is to have received the DPA notice. One large provider has at least set a clock: Microsoft's May 2026 DPA update gives 30 days' notice before new AI subprocessors (down from six months), with a customer right to disable the subprocessor for a period after notice.
- **Notice periods are bimodal by provider — and this may be the most consequential stat in the report.** Anthropic's median notice on model retirement is **62 days** (minimum 60). OpenAI's median across all obligation-triggering events is **182.5 days**. A vendor building on Anthropic gets roughly one-third the migration runway, and a 60-day floor sits *below* the change-control window many health-tech compliance programs run on.

2. The buyer side has already written the spec

HSCC model contract language (April 2026). The Health Sector Coordinating Council's Third-Party AI Risk and Supply Chain Transparency Guide includes sample commercial contract language (Appendix E) and sample BAA language (Appendix F). These are recommendations, not signed contracts — but they are the health sector's own template, and they are specific:

OBLIGATION	WINDOW	MODE
Planned update — release notes + validation evidence	30 days	Advance
Emergency security patch	72 hours	Advance
Updates modifying model architecture or data processing	Customer approval before deployment	Consent gate
Change to third-party AI services processing PHI	Contract amendment + privacy impact assessment + governance approval, before implementation	Consent gate
Performance degradation beyond threshold	48 hours	Retrospective
Security incident / model failure / anomaly	2 hours (Critical) / 24 hours	Retrospective
AI-specific PHI breach	2 days from discovery	Retrospective
Audit request	30 days' notice	On request

Bracketing note: HSCC's model text uses square brackets for values the parties are meant to fill in. Three windows above appear bracketed in the clause text — the 48-hour degradation notice, the 30-day audit notice, and the 2-day AI-breach window (though the 2-day figure appears unbracketed in the guide's body) — meaning HSCC prompts a number there without recommending one. The 30-day advance, 72-hour patch, and 2-hour/24-hour incident figures are unbracketed in the example clause language.

Read the fourth row again. Under the health sector's model language, a vendor switching upstream AI providers isn't sending a notification — it's reopening the contract. And the model BAA language (Appendix F) names the upstreams explicitly: it requires BAAs with "AWS, Azure, Google Cloud AI services, OpenAI, Anthropic, etc." before PHI is processed through them.

Put sections 1 and 2 side by side and the tension is the story: the model contract language asks for **30 days' advance notice** of model changes, from vendors whose own suppliers give them a **60-day floor**, no announcement dates (Azure, Google), and **no subprocessor changelog at all**. The contract asks for something the supply chain barely permits — unless the vendor can generate the evidence itself, at runtime.

It's not just guidance. In-force contracts exist. Medica — a Minnesota-based payer — publishes an Artificial Intelligence Addendum (v.2026.01) whose posture is stronger than any notice clause: §2.1 requires disclosure of model name, use case, and data used, plus *prior written consent* before any AI touches the engagement (revocable at Medica's discretion); §2.2 requires the vendor to **cease using any new or previously undisclosed AI until approved in writing**. A model swap under that contract is a stop-work event. §2.12 grants audit access to "training data, algorithms, audit trails, logs and other materials that document the inputs and outputs of Vendor's AI process" — at no cost, on request. The State of California's GenAI Special Provisions (January 2025) independently require contractor disclosure of training-data quality metrics, output quality metrics, and "reasonable access to the GenAI security logs, latency statistics, and other related GenAI security data."

One precision note we insist on because it is the easiest place to blur: **HSCC's 30 days is advance; Medica's 30 days (\$2.5, change notification) is retrospective — within 30 days after a modification.** Two 30-day clocks that mean opposite things. Any vendor or buyer merging them is mis-stating one side.

The newer May 2026 HSCC guide raises the bar again. The AI Cyber Governance Framework Implementation Guide — which describes itself as complementing the April guide, to "be used in conjunction with" it — requires post-update validation evidence ("updated model cards, performance benchmarks, and bias testing results") *as a condition of update acceptance*, and for agentic AI states that "every action... must be logged with sufficient detail to support forensic reconstruction: what action was taken, what data was accessed or modified... Logs must be tamper-resistant." Its vendor-incident section sets tiered SLAs in unbracketed body text — notification within 24 hours of vendor awareness for critical incidents (patient safety or PHI exposure) with assessment within 48 hours, 5 business days for high-risk, 10 for medium, and 30 days for vulnerability disclosures including supply-chain dependencies — and its subprocessor section asks vendors to give advance notice of material subprocessor changes "with opportunity for the healthcare organization to object or terminate."

The change triggers finally have a definition. CHAI's Applied Model Card is, as far as we can find, the only published source that defines a material AI change with specificity: a base foundation-model swap (e.g., GPT-4 → Llama), a new training dataset with a new patient population, a new imaging acquisition type, or a major version bump — each requiring a revised model card tied to the software revision number. Every contract template we surveyed leaves "material change" undefined. (Also worth knowing: 36 partner organizations back the Applied Model Card initiative; as of this writing, one published card is publicly visible on the registry.)

3. The supply side: almost nothing

We looked for the instrument that would satisfy these obligations. Findings:

- **SEC EDGAR full-text search (coverage 2001-present) contains zero filed exhibits matching "artificial intelligence addendum"** — verified against a live control query, as of August 3, 2026. Variants ("artificial intelligence amendment," "AI rider," "artificial intelligence exhibit," "generative AI addendum") also return zero. ("AI addendum" returns nine hits: one 2006 filing where "AI" abbreviates Allsports International Ltd., and eight copies of the Delaware Life/Zinnia filing discussed below — none a standalone AI addendum instrument.)
- **The obligations are appearing — buried, unlabeled, inside master agreements.** The closest thing to a counterexample: a services agreement filed June 30, 2026 (Delaware Life / Zinnia Tech Solutions) whose Section 22 is captioned "Artificial Intelligence," with a training-consent gate and a dedicated AI compliance exhibit incorporated by reference. The substance is arriving. The reusable, standalone instrument is not.
- **Neutral and vendor-side templates are nearly silent.** Common Paper's AI Addendum (CC BY) and Bonterms' AI Standard Clauses (CC0) contain no model-change notification duty. A healthcare EHR vendor's own click-through AI addendum contains no change notification, no audit rights, and no validation duty. Across every vendor-side and neutral instrument we reviewed, only two fixed advance-notice windows for an AI change exist as of this writing: Google Cloud's 30 days before re-designating a feature as "Advanced AI," and Microsoft's May 2026 DPA update — 30 days' notice before engaging a new AI subprocessor (down from six months), paired with a time-limited customer right to disable it.
- **Standard health-tech BAAs don't cover it.** In the five BAA exhibits we pulled from EDGAR filings, subprocessor terms are flow-down rather than notification (the customer has no right to be told who was added), audit rights mostly run to the HHS Secretary rather than the customer, and not one states a hard incident-notice day-count in its public text.

4. What we're doing about it

Today we're publishing the **AI Change Notice** schema (v0.1): a one-file, machine-readable format for the notice the guidance already describes — an operationalization of the sample contract language in HSCC's public guidance (Appendices E/F), with attribution, rather than an invention. v0.1 deliberately ships narrow: change-event types and clause categories are free-form fields, because no authoritative taxonomy has been adopted yet. A controlled vocabulary — including alignment with the material-change triggers defined in CHAI's Applied Model Card — is scoped for v0.2, once upstream licensing and taxonomy questions settle. Spec, validator, and worked example at github.com/Shaachi1/ai-change-notice. Contributions and criticism welcome; we commit to donating the schema to a neutral standards body.

What we want from readers, in order: hospitals and TPRM teams who would *accept* a notice in this format (the readers matter more than the writers); health-AI vendors who want to publish their first notice; and corrections — every claim above

carries its source, and we will publish any correction with the same prominence as the original.

Methodology and honesty notes

Built from primary sources reviewed July 30 – August 3, 2026: the HSCC April and May 2026 guides including Appendices E–G of the April guide; the Joint Commission × CHAI RUAIH guidance; CHAI Applied Model Card v0.1 documentation and registry; the Medica AI Addendum v.2026.01; California's GenAI Special Provisions; Common Paper, Bonterms, Google Cloud, SmartRecruiters, and Enterprise Health/MIE published terms; SEC EDGAR full-text search with control queries; and the official model-lifecycle pages of OpenAI, Anthropic, Azure OpenAI, and Google Vertex AI. Scope limits stated inline: EDGAR full-text coverage begins 2001; HSCC appendix text is model language, not signed contract language; an earlier internal estimate of "one platform change every two weeks" was corrected to the defensible "every three weeks" after the event feed was verified against official sources — the corrected number is the one printed here.

Gatekeeper is the evidence layer for healthcare AI: per-call PHI redaction before prompts leave the vendor's tenant, prompt-injection screening, and an audit-evidence packet mapped to SOC 2/HIPAA controls. That position is why we care about runtime evidence; readers should weigh our incentives accordingly.

© 2026 K2 Corp (d/b/a Gatekeeper). Corrections: this report will publish any correction with the same prominence as the original claim. Contact: dilip@slaysales.com · AI Change Notice schema: github.com/Shaachi1/ai-change-notice