

EVEGLYPHDESIGN · REFERENCE MODEL

# Global Compliance Framework Assessment

---

Document ID EgD-GCF-001 · Key ID EgD-KEY-2026-07 · Status reference model, v1.0 · Issued 2026-08-02T02:43:12Z

SHA-256 of source cac3093ee4f44e3ad1f840dfe93090970faf37c1c569edf92b088e43101b283a

A buyer-side reference model for enterprise clients evaluating an AI proposal from an incumbent vendor, scored on data sovereignty rights and access. Companion instrument to the seven-point enterprise-grade audit published in this repository.

---

## Why this exists

The dangerous AI proposal is not the one from a startup. It is the one that arrives as a three-page amendment to an existing master agreement with an incumbent the enterprise already runs its business on — the ERP vendor, the CRM vendor, the productivity suite, the core banking platform. It is priced as an uplift, scoped as a feature, and reviewed as a renewal. It is almost never reviewed as what it actually is: a change in who holds title to the enterprise's operational meaning.

Three properties make the incumbent case distinct.

- The switching cost is already sunk. The buyer has no credible walk-away, so the ordinary commercial discipline that governs a competitive selection does not operate.
- The data is already resident. The vendor is not asking for access — it already has it. The amendment asks only to change the purpose of processing, which is the single most consequential and least negotiated clause in the document.
- The derived artifacts are new. Embeddings, indexes, fine-tunes, evaluation sets, and agent traces did not exist when the original agreement was signed. Silence in the original contract is routinely read by the vendor as permission.

This assessment gives the buyer a fixed, jurisdiction-anchored set of questions, an evidence standard for each, and a scoring sheet that produces a defensible recommendation rather than a feeling. It is written to be handed to a vendor unchanged.

**Sovereignty is not where the data sits. Sovereignty is who can still operate the business on the day the vendor relationship ends badly.**

## Scope and how to score

Assess a single named proposal, not a vendor in the abstract. Each of the eight control domains is scored 0, 1, or 2:

- 2 — Evidenced. The vendor has produced a contractual commitment or an artifact that can be tested. Documentation the buyer can hold after termination.
- 1 — Asserted. The vendor states the position in a proposal, trust centre, or sales call, but it is not in the contract and cannot be tested.
- 0 — Absent, refused, or reserved. Including any answer of the form "that is covered by our standard terms" where the standard terms do not cover it.

Marketing collateral, trust-centre pages, and security whitepapers are evidence of assertion only. They are unilaterally amendable by the vendor and score a maximum of 1. A certification scope statement, an executed data processing agreement, an order form clause, or a working export is evidence and scores 2.

## Part I — The eight sovereignty control domains

### SD-1 · Title to derived data

The question. Who owns the embeddings, vector indexes, fine-tuned weights, prompt libraries, evaluation sets, and agent execution traces generated from the enterprise's data — and does that ownership survive termination?

Why it decides the deal. Input data ownership is almost always conceded to the customer and is almost always irrelevant. The operating value migrates into the derived layer. A vendor that concedes input ownership while retaining the index has taken the asset and left the buyer the receipt.

Evidence required. An express contractual grant of customer title, or an irrevocable, perpetual, transferable licence, over each named derived artifact class; plus written confirmation of which artifacts are exportable in a documented format.

Scoring. 2 — title or perpetual transferable licence over all named classes, exportable. 1 — customer owns inputs, derived artifacts unaddressed. 0 — vendor asserts ownership of derived artifacts, or the contract is silent and the vendor declines to clarify in writing.

## SD-2 · Purpose limitation and model-improvement use

The question. Can enterprise content, prompts, outputs, telemetry, or support-ticket data be used to train, tune, evaluate, or benchmark any model that serves another customer — by the vendor, by a subprocessor, or by a foundation-model provider upstream?

Anchor. [GDPR Article 5\(1\)\(b\) purpose limitation](#) and [Article 28\(3\)\(a\)](#), which require a processor to act only on documented instructions. Under the EU AI Act, providers of general-purpose AI models have carried disclosure and documentation obligations [since 2 August 2025](#).

Evidence required. A default-off, contractually binding no-training commitment covering the full chain including upstream model providers; the same commitment applied to abuse-monitoring and human-review pathways, with retention periods named in days.

Scoring. 2 — binding, chain-wide, default off, retention stated. 1 — opt-out exists but is default-on, or covers the vendor but not upstream providers. 0 — training rights retained, or "aggregated and de-identified" carve-outs left undefined.

## SD-3 · Extractability — the Full-Text Timestamped Extract standard

The question. Can the enterprise obtain, on demand and without vendor assistance, the complete text of every interaction, document, prompt, output, and decision record the platform holds on its behalf — as full text, timestamped, in a documented, non-proprietary format?

Anchor. [GDPR Article 20 data portability](#) for personal data, extended here by EVE Glyph Design canon to the whole enterprise record via the Full-Text Timestamped Extract requirement in the data-liberation posture.

Evidence required. A live, buyer-executed export producing full text with timestamps and stable identifiers — not a screenshot of an export button, and not a summary or sampled export. A published schema. A documented API with rate limits that permit a full extract inside the notice period.

Scoring. 2 — buyer has personally executed a complete export and validated the schema. 1 — export documented but untested, or partial. 0 — no bulk export, PDF-only export, or extract available solely as a professional-services engagement.

## SD-4 · Switching and functional equivalence

The question. Can the enterprise move to another provider or to its own infrastructure and reach functional equivalence, and what does the vendor charge to let it happen?

Anchor. The [EU Data Act \(Regulation 2023/2854\)](#), applicable [from 12 September 2025](#), requires providers of data processing services to remove pre-commercial, commercial, technical, contractual and organisational obstacles to switching, and to contract for a defined notice period and transition window. Switching charges are capped at directly incurred cost now and are [prohibited outright from 12 January 2027](#).

Evidence required. Contractual switching-assistance obligations, the maximum transition period in days, the fee schedule, and a written statement of the vendor's Data Act position. A vendor still quoting egress or switching fees in a proposal that runs past January 2027 has either not read the regulation or is pricing the buyer's ignorance.

Scoring. 2 — switching clause present, assistance obligations named, no charge beyond direct cost, Data Act position stated. 1 — assistance offered at time-and-materials, no committed window. 0 — no switching provision, or egress and switching fees priced as revenue.

### SD-5 · Residency, transfer, and lawful-access exposure

The question. In which jurisdictions is data stored, processed, and inferenced — and which governments can lawfully compel disclosure from the vendor or its parent, regardless of where the bytes sit?

Why residency alone is not an answer. Inference frequently crosses the boundary the storage contract respects. A vendor may hold data in-region and route the prompt to a model endpoint elsewhere. Ask about the inference path explicitly, and about the model provider's own logging.

Anchor. GDPR Chapter V transfer rules and the Schrems II adequacy and supplementary- measures analysis; the US CLOUD Act extraterritorial production power over providers subject to US jurisdiction; and national regimes including [India's DPDP Act and DPDP Rules 2025, in force on staggered dates through 13 May 2027](#), Brazil's LGPD, Canada's PIPEDA and Quebec Law 25, and China's PIPL.

Evidence required. A jurisdiction map covering storage, processing, inference, telemetry, and support access; the transfer mechanism relied on for each; the group corporate structure that determines lawful-access exposure; and the vendor's law enforcement request transparency report and notification commitment.

Scoring. 2 — full map including inference and support tiers, mechanisms named, parent structure disclosed. 1 — storage residency only. 0 — "global infrastructure", or refusal to name the inference jurisdiction.

### SD-6 · Substrate and subprocessor transparency

The question. Which foundation models, in which versions, from which providers, sit behind the surface the vendor is selling — and how is the buyer notified when they change?

Why it belongs in a compliance framework. A model rotation is a subprocessor change and a risk-posture change at the same time. It alters residency, training exposure, failure modes, and the applicable regulatory classification, and it is routinely made without notice because the contract treats it as an implementation detail.

Anchor. [GDPR Article 28\(2\) and 28\(4\)](#) on authorised subprocessors, read together with the enterprise-grade audit's roster requirement in this repository.

Evidence required. A published, versioned foundation-model roster; a subprocessor list with change notification and an objection right; substrate-selection controls that let the buyer pin or exclude specific models.

Scoring. 2 — versioned roster, notification with objection right, selection controls available. 1 — subprocessor list published, models unnamed. 0 — model providers treated as confidential.

### SD-7 · Access, logging, and audit rights

The question. Can the enterprise see, export, and independently audit who accessed what — including vendor personnel access — and can it obtain the agent decision logs needed to answer a regulator?

Anchor. [EU AI Act Article 12 logging and Article 26 deployer obligations](#); [DORA Article 30](#), which for financial entities requires full access, inspection, and audit rights in contracts covering critical or important functions; NIS2 supply-chain security duties; and [ISO/IEC 42001](#) as the AI-management-system baseline a mature vendor should already hold.

Evidence required. Customer-exportable audit logs with a stated retention period; vendor-personnel access logging surfaced to the customer; a right to audit or a recognised third-party report whose scope statement actually covers the AI service rather than the legacy platform. Check the scope statement — an incumbent's long-standing SOC 2 or ISO 27001 certificate frequently excludes the new AI service entirely.

Scoring. 2 — exportable logs including vendor access, audit right or in-scope attestation. 1 — logs visible in-console only, or attestation with unverified scope. 0 — no customer-accessible logging of vendor access.

### SD-8 · Exit, deletion, and post-termination survivability

The question. On the day the relationship ends, what does the enterprise still hold, what does the vendor still hold, and can the business keep running?

Evidence required. A deletion certificate covering derived artifacts and backups with a stated timeline; a post-termination data-availability window; confirmation that fine-tuned weights derived from customer data are deleted or delivered; and a survivability statement naming which business processes stop working at termination.

Scoring. 2 — certificate covering derived artifacts, window stated, survivability documented. 1 — deletion of primary data only. 0 — no deletion commitment for derived artifacts, or immediate cut-off with no export window.

---

## Part II — Jurisdictional register

The dates below are the live compliance anchors as at 1 August 2026. Confirm before relying on them in a specific transaction.

- EU AI Act (Regulation 2024/1689) — prohibitions and AI-literacy duties applied from 2 February 2025; general-purpose AI model obligations and governance from 2 August 2025; [Article 50 transparency obligations apply from 2 August 2026](#). Under the Digital Omnibus adopted in mid-2026, high-risk obligations for stand-alone Annex III systems were [deferred to 2 December 2027, and for AI embedded in regulated products to 2 August 2028](#). Buyer note: the deferral is not a reprieve. It moves the vendor's deadline, not the buyer's exposure — the deployer obligations still land on the enterprise, and a vendor citing the omnibus as a reason to defer disclosure is transferring its own risk to the customer.
- EU Data Act (Regulation 2023/2854) — applicable 12 September 2025; switching obstacles must be removed; switching charges [prohibited from 12 January 2027](#).
- GDPR (Regulation 2016/679) — [Article 20](#) portability, [Article 28](#) processor and subprocessor duties, Chapter V transfers, Article 22 automated decision-making.
- DORA (Regulation 2022/2554) — applicable from 17 January 2025 to EU financial entities; [Article 30](#) mandatory contractual provisions and the register of information for ICT third parties.
- NIS2 (Directive 2022/2555) — supply-chain security and incident reporting duties transposed nationally; relevant wherever the AI service touches essential or important entity operations.
- United States — CLOUD Act extraterritorial production; sectoral regimes (HIPAA, GLBA, FERPA); state comprehensive privacy statutes with automated-decision and profiling provisions; NIST AI Risk Management Framework as the common voluntary control language.
- Canada — PIPEDA federally and Quebec Law 25, including automated-decision disclosure and the right to information about the decision logic.
- India — [DPDP Act 2023 and DPDP Rules 2025, notified November 2025 with obligations phasing to 13 May 2027](#).
- Brazil — LGPD, including the right to review of automated decisions.
- China — PIPL, plus the generative-AI measures and security-assessment requirements governing cross-border transfer.
- Standards baseline — [ISO/IEC 42001](#) AI management systems, ISO/IEC 27001 information security, and SOC 2 Type II. Always read the scope statement, never the badge.

---

## Part III — The four veto conditions

Any one of these is an unconditional fail regardless of total score. They are not weightings. They are the conditions under which the enterprise's own record ceases to be its own.

1. Derived-artifact capture. The vendor asserts ownership of, or refuses to license, embeddings, indexes, or fine-tunes built from enterprise data.
2. Unbounded training rights. Enterprise content may be used to improve models serving other customers, and the right cannot be contractually switched off end to end.

3. No independent extraction. There is no path to a complete, timestamped, full-text export that the buyer can execute without vendor cooperation.
4. Undisclosed inference jurisdiction. The vendor will not state where inference occurs or which governments can compel disclosure from it or its parent.

---

## Part IV — Scoring sheet

Maximum 16 points across eight domains.

- 14–16, no veto — Proceed. Sovereignty posture is contractually defensible.
- 10–13, no veto — Proceed with conditions. Convert every domain scored 1 into a contractual commitment before signature. Asserted is not evidenced.
- 6–9 — Renegotiate. The proposal is a feature purchase carrying a governance liability. Do not sign on the vendor's paper.
- 0–5, or any veto condition — Decline or restructure. Restrict to a bounded pilot on non-critical, non-personal data with an explicit termination date and no derived-artifact retention.

Record each score with the evidence artifact that justifies it, and file the completed assessment with the contract. The scoring sheet is the audit trail; a score without a named artifact is an opinion.

---

## Part V — Contract language hooks

Where a domain scores below 2, the remedy is a clause, not a follow-up call. Insert:

1. Derived data title — customer owns, or holds a perpetual irrevocable transferable licence to, all named derived artifact classes, enumerated rather than described.
2. No training without instruction — default off, chain-wide including upstream model providers, with an audit right against it.
3. Extraction warranty — complete, timestamped, full-text export in a documented format, available throughout the term and for a stated window after termination, without professional-services fees.
4. Switching assistance — obligations, maximum transition period in days, and no charge beyond directly incurred cost, expressly referencing the Data Act position.
5. Substrate change notice — advance written notice of foundation-model or subprocessor change, with an objection right and a termination trigger.
6. Jurisdiction warranty — storage, processing, inference, telemetry, and support jurisdictions warranted, with notice of change.

7. Audit and log access — exportable logs including vendor-personnel access, with a retention floor and an in-scope third-party attestation.
8. Exit and deletion — deletion certificate covering derived artifacts and backups, on a stated timeline, with a survivability statement.

---

## Part VI — Ten patterns to flag in an incumbent proposal

1. The AI capability arrives as an amendment rather than a new agreement.
2. The purpose-of-processing clause is broadened while the data-protection exhibit is left untouched.
3. "Aggregated and de-identified" appears without a definition or a re-identification standard.
4. The existing certification is cited without a scope statement covering the AI service.
5. Foundation-model providers are named as confidential.
6. Export is available, but only as summaries, reports, or PDFs.
7. Egress or switching fees are quoted for a term running past January 2027.
8. Inference location is described as "global" or "regional" without naming a country.
9. Deletion commitments cover primary records and are silent on indexes and fine-tunes.
10. The commercial uplift is bundled into a renewal so the AI terms are never separately reviewed.

---

## Crosswalk to the seven-point enterprise-grade audit

This assessment is the buyer-side instrument; the [seven-point audit](#) is the vendor-side standard. They meet at three points. Audit point 1, the foundation-model roster, is the evidence that satisfies SD-6. Audit point 4, substrate-selection controls, is the control that raises SD-6 from 1 to 2. Audit point 5, liability commensurate with the tier claim, is what makes the Part V clauses enforceable rather than decorative. A vendor that passes the seven-point audit will pass most of this assessment without negotiation. That is the point of publishing both.

The framework also carries the EVE Glyph Design mirror posture: the correct architecture for an AI layer over an incumbent estate is one that mirrors and never cannibalises — read-only against the system of record, no displacement of the authoritative ledger, and full portability of the data stores it builds. A proposal that requires the system of record to move in order to make the AI work has failed SD-4 before it has been scored.

---

## Use and limits

This is a procurement and governance instrument, not legal advice. Regulatory dates move — the EU AI Act high-risk deferral in 2026 is the recent proof. Verify each anchor against the primary source for the jurisdictions that bind the specific transaction, and have counsel review any clause taken from Part V before it is inserted.

Reuse is permitted and encouraged, including by vendors preparing to answer it.

---

## Version and provenance

- Document ID EgD-GCF-001 · Key ID EgD-KEY-2026-07 · Status reference model, v1.0
- Source of record — the Markdown source in the [enterprise-grade-criteria repository](#). The controlled PDF carries a SHA-256 of that source in its footer; a PDF whose hash does not match the committed source is not a controlled copy.
- Revision policy — the jurisdictional register in Part II is the volatile section and is revised when an anchor date moves. The eight control domains and the four veto conditions are stable and change only by version increment.
- Related canon — the seven-point enterprise-grade audit in this repository; the [data-liberation-kit](#) for the Full-Text Timestamped Extract standard referenced in SD-3; the [truth-ledger](#) for the filed lock-in record the control domains are drawn from — LOCKIN-01 maps to SD-4, PORTABILITY-06 to SD-3, GATING-03 to SD-7.
- Routing — scoring a live vendor proposal with this instrument is Class F.2 in the [task routing index](#). Filing a defect against a tier claim already observed remains Class F.1.