CoreSR Inc.
c/o LaBarge Weinstein LLP, 800-515 Legget Drive, Kanata, ON K2K 3G4, Canada
help@coresr.ai · coresr.ai

PUB-006 · Version 1.4
Effective: 2026-11-17
Public Procurement & Governance

AI Governance, Data Provenance & Zero-Egress Architecture

Controls Covered: CoreSR AI Governance Controls 4.1 – 4.6
Security, Privacy & Legal Contacts: security@coresr.ai · legal@coresr.ai

Hospitals, academic medical centres, clinical guideline societies, and health technology assessment (HTA) agencies evaluate AI systems against questions that generic cloud security frameworks do not answer: *What were your models trained on? Is our unpublished systematic review or clinical guideline used to train your models? Does our literature corpus leave your infrastructure for commercial chatbot APIs? Can our authors prove which exact model version produced a screening or extraction proposal when peer reviewers ask?*

Because CoreSR trains and serves its own domain-specific biomedical models (**TITAN-SR**, **GLASS**, and **ROB BERT**) rather than wrapping third-party consumer LLM APIs, CoreSR enforces verifiable architectural controls across all six areas below.

Control 4.1 — Training Data Provenance & Licensing Basis

CoreSR operates three specialist model families in production. None of CoreSR's models is trained on patient medical records (PHI), private customer manuscripts, or unauthorized shadow libraries.

1. TITAN-SR (`v6-2026-10-15` GPU & `dual3-2026-07-12` Free/Trial CPU)

- **Role:** Calibrated title-and-abstract systematic-review eligibility screening.
- **Training Corpus:** 3,097,955 bibliographic citation records (315,052 final-inclusion positives) across 19,926 systematic-review groups (snapshot October 14, 2026), with 80 held-out calibration reviews and 60 unseen test reviews (91,538 records) strictly isolated from training.
- **Data Sources & Licensing Basis:** Publicly distributed scholarly bibliographic metadata — titles, abstracts, publication types, and indexing descriptors retrieved through authorized public scholarly indexes (**PubMed / NCBI NLM**, **Europe PMC**, **OpenAlex** CC0, **Crossref** public metadata, and **ClinicalTrials.gov**) paired with factual citation inclusion outcomes published in peer-reviewed systematic reviews, plus de-identified screening labels where permitted under Individual/Free account terms (§4.2). Underlying encoder architectures derive from permissively licensed (**Apache-2.0 / MIT**) biomedical transformer checkpoints.

2. GLASS (`v3` — Grounded Literature Appraisal and Synthesis System)

- **Role:** Evidence-grounded full-text screening, structured PICO/outcome extraction, Risk-of-Bias signalling-question appraisal, GRADE certainty rationale drafting, and synthesis assistance across six task-specialized routes (65,536-token context window).
- **Training Corpus & Licensing Basis:**

1. **Foundation Weights:** Built upon permissively licensed (**Apache-2.0**) open-weight foundation models permitting commercial adaptation, self-hosting, and full weight custody without external vendor callbacks.
2. **Biomedical Domain Pre-Adaptation:** Continued domain adaptation on the 3,097,955-record TITAN bibliographic screening corpus.
3. **Supervised Fine-Tuning & Alignment:** Curated methodology instruction sets, open-access biomedical full-text articles (PMC Open Access Subset / CC-BY / CC0), published Risk-of-Bias and GRADE methodological rubrics, and internally verified structured appraisal worksheets. GLASS is trained to perform structured methodological reasoning and verbatim source-quote grounding — never to memorize or reproduce paywalled publisher full texts.

3. ROB BERT (`v6.1`)

- **Role:** Specialist two-stage Cochrane RoB 2 signalling-question reader and 22-head domain checker across all five RoB 2 bias domains (8,192-token context, calibrated confidence gating).
- **Training Corpus & Licensing Basis:** Trained on 4,237 deduplicated structured RoB 2 worksheets and evaluated on 532 held-out test worksheets paired with open-access trial passages under permissive encoder licences.

Control 4.2 — Customer Content & Model Training (Single Unambiguous Position)

CoreSR applies a **single, uniform contractual and technical policy** across its Terms of Service (https://coresr.ai/terms), Privacy Policy (https://coresr.ai/privacy), PIPEDA & GDPR Data Processing Addendum (https://coresr.ai/api/public/dpa), and institutional procurement commitments:

1. **Team & Institution Workspaces — Never Used for Model Training:**
   - Customer Content in **Team** and **Institution** workspaces (protocols, search strategies, uploaded PDFs, screening decisions, extractions, Risk-of-Bias judgements, GRADE tables, meta-analytic datasets, and manuscript drafts) is **never** used to train, fine-tune, distil, or evaluate CoreSR models.
   - **Technical Enforcement:** Workspace-tier exclusions are enforced directly in server-side export pipelines at query time so every project and member belonging to a Team or Institution workspace is automatically excluded.
   - **Bilateral Research / Co-Development Carve-Out (Explicit Written Agreement Only):** Where an academic consortium, guideline society, or institutional research partner wishes to conduct joint benchmark evaluation or bespoke adapter co-development on *its own* designated corpus, that collaboration occurs **exclusively** under a separate, bilaterally signed written Research / Co-Development Agreement specifying dataset scope, governance, and intellectual property — never under standard subscription terms.
2. **Individual & Free Accounts — Transparent De-Identification & One-Click Opt-Out:**
   - Standalone Individual and Free accounts may enable **"Exclude my reviews from model training"** in one click under **Account Settings → Privacy**, taking effect immediately and prospectively.
   - Even when an Individual/Free user leaves the default setting active, **manuscript drafts, private reviewer notes, unpublished proprietary documents, and personal identifiers are never exported** — only de-identified bibliographic title/abstract eligibility pairs stripped of user and project identifiers.

Control 4.3 — Prompt & Inference Telemetry Retention

Storage Layer	Recorded Metadata	Contains Customer Prompts or Study Text?	Retention Schedule
Daily Quota Metering	Per-account daily operation count	**No** — integer counter only	Active account lifetime (purged upon account deletion)
Operational Cost & Latency Telemetry	Operation category, model identifier, input/output token counts, latency, HTTP status, estimated compute cost	**No** — schema excludes prompts, responses, URLs, and study text	365 days rolling operational analytics
Self-Hosted GPU Inference Workers (CoreSR & RunPod Standby)	Ephemeral in-memory tensor execution and aggregate health metrics	**No persistent prompt logging** — context exists in volatile memory solely during request execution	**0 seconds** post-response
Customer Review Work Product	Structured decision, rationale, verbatim study quote, confidence, and pinned model version within the Customer's project	**Yes — as Customer-owned project data** visible in the Customer's workspace and exports	Controlled by Customer; purged from live storage following self-service account deletion
Application Error Diagnostics	Sanitized error status codes and service health diagnostics	**No** — prompt bodies and study text are excluded from diagnostic logs	30 days rolling on encrypted storage

Control 4.4 — Model Version Pinning & Scientific Reproducibility

Methodological reproducibility requires knowing **which exact model artifact** generated an AI proposal so systematic reviews remain auditable for journal peer review, living-guideline updates, and regulatory/HTA inspection:

1. **Per-Output Artifact Pinning:** Every AI-assisted screening score, extraction item, Risk-of-Bias proposal, and GRADE certainty suggestion persists its exact model family, release tag, operating-point threshold, and verified artifact manifest digest alongside the record.
2. **Batch Generation Isolation:** Server validation verifies model artifact cryptographic digests before accepting screening scores and prevents mixing incompatible calibration generations within a single screening batch unless the reviewer explicitly initiates reclassification.
3. **Preserved in Customer Exports:** Both the **Screening Audit CSV** and the **One-Click Full Account ZIP Export** include pinned model release versions, sensitivity targets, applied thresholds, and human override indicators.
4. **Dated Public Release Log:** Every production model promotion, corpus snapshot, and calibration update is documented in the CoreSR Model Version Changelog (https://coresr.ai/api/public/model-changelog).

Control 4.5 — Human-in-the-Loop as an Architectural Invariant

In CoreSR, *"the platform proposes; a qualified human reviewer decides"* is enforced at the **data-model and workflow state-machine level**:

1. **Non-Destructive Dual-State Records:** AI screening classifications and reviewer adjudications are stored in separate fields; full-text screening preserves independent dual-reviewer and tiebreaker states without overwriting human decisions.

2. **Mandatory Verbatim Source Grounding:** Every GLASS-extracted study value must carry a locatable **verbatim PDF quote and page reference** — unlocatable quotes are rejected automatically — and remains flagged unverified until a human reviewer inspects the highlighted PDF passage and approves or edits the value.
3. **Deterministic Methodological & Statistical Rules:** Published Risk-of-Bias domain decision trees (such as Cochrane RoB 2) and quantitative CoreGRADE domains (**Risk of Bias**, **Imprecision**, and **Publication Bias**) execute **100% deterministically in code** from study data and statistical thresholds rather than generative prose estimation. Reviewer-approved judgements are never overwritten on subsequent runs.

Control 4.6 — Third-Party Commercial LLM Egress

Unlike wrappers built atop external consumer chatbot APIs, CoreSR executes screening and evidence synthesis on **its own self-hosted model weights**:

1. **Server-Enforced Vendor Egress Block (Team & Institution):** Every request originating from a **Team** or **Institution** workspace executes inside a mandatory server-enforced zero-vendor-egress policy scope. Any code path attempting to invoke an external commercial LLM API provider (including OpenAI, Anthropic, Google Gemini, or DeepSeek) fails closed in-process before any outbound network request can open, with the single exception below.
    - **The exception — Kaplan–Meier figure reading (KM Cor), effective November 17, 2026:** When a reviewer runs **KM Cor** on a Kaplan–Meier figure, the cropped image of that figure, its caption, and the review's subgroup names are sent to **Google's Gemini vision model** through a commercial model API, in Team and Institution workspaces as on every other plan. Nothing else is sent: not the PDF, not any other page, and no screening, extraction, risk-of-bias or GRADE data. KM Cor runs only when a reviewer starts it on a figure, and a figure can instead be traced by hand, which uses no AI. CoreSR's own vision model does not yet read Kaplan–Meier curves as accurately as the external model, so the figure read uses the external model on every plan until it does. The exception is one call in server code and is covered by automated tests that confirm every other call from these workspaces is still refused.
2. **Controlled Execution Topology:**
    - **European Union Primary Production (Helsinki, Finland):** Hosts application orchestration, encrypted relational storage, and dedicated CPU screening services.
    - **Canadian Dedicated GPU Infrastructure (Ontario, Canada):** Reached exclusively over an authenticated encrypted point-to-point tunnel; serves primary TITAN-SR, GLASS, and ROB BERT inference on CoreSR-operated GPUs.
    - **Standby GPU Containers (RunPod, Inc. — SOC 2 Type II):** Used only when CoreSR's private servers are down or unavailable. Requests then route over TLS 1.3 to private, stateless containers executing **CoreSR's own GLASS weights** in volatile GPU memory with zero persistent disk, zero prompt logging, and zero third-party model API calls.

ISSUED ON BEHALF OF CORESR INC.

Tyler Pitre, MD PhD FRCPC
Founder and Chief Executive Officer, CoreSR Inc.
tyler.pitre@coresr.ai · coresr.ai