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

PUB-010 · Version 1.3
Effective: 2026-11-17
Public Procurement & Governance

Institutional Security, Privacy & AI Questionnaire Response Pack

Standardized Procurement Responses for Hospitals, Universities, Guideline Societies & HTA Agencies
Security, Privacy & Legal Contacts: security@coresr.ai · legal@coresr.ai

Every response below is reconciled against CoreSR's live production controls, Privacy Policy (https://coresr.ai/privacy), Terms of Service (https://coresr.ai/terms), PIPEDA & GDPR Article 28 Data Processing Addendum (https://coresr.ai/api/public/dpa), Security Controls Statement (https://coresr.ai/api/public/controls-statement), and AI Governance Specification (https://coresr.ai/api/public/ai-governance).

Section A — Governance, Certifications & Independent Testing

Q1. Does CoreSR hold ISO 27001 or SOC 2 certification?
Answer: CoreSR does not hold an external ISO 27001 or SOC 2 certificate today. CoreSR operates a formal internal Information Security Management System (ISMS) modelled on **ISO/IEC 27001:2022**, comprising signed security and continuity policies, a complete ISO/IEC 27001:2022 Statement of Applicability continuously verified against live infrastructure (zero control drift), a formal Risk Register, quarterly privileged-access reviews, and continuous automated posture verification. CoreSR is pursuing **SOC 2 Type II** directly (rather than a single-day Type I snapshot), with the Type II observation window scheduled to open in **Q1 2027** (target attestation report **Q4 2027**).

Q2. Has CoreSR undergone penetration testing and vulnerability scanning?
Answer: CoreSR runs continuous automated static, dependency, secret, and dynamic security testing: (1) blocking CI/CD static application security testing (SAST), software composition analysis (SCA), and secret detection on every pull request and weekly schedule; (2) automated weekly dynamic application security testing (DAST) scans against production; and (3) a public safe-harbor Coordinated Vulnerability Disclosure programme (https://coresr.ai/.well-known/security.txt). Validated vulnerabilities are governed by binding remediation SLAs (**Critical ≤ 7 days, High ≤ 30 days, Medium ≤ 90 days**). CoreSR commissions an independent accredited third-party penetration test upon institutional purchase requirement or at SOC 2 Type II kickoff (Q1 2027), whichever occurs first.

Section B — Data Scope, Privacy, Hosting & Cross-Border Residency

Q3. What categories of data does CoreSR store or process? Does CoreSR process PHI?
Answer: CoreSR processes **published scientific research literature and methodological review work product only** (search strategies, citations, uploaded study PDFs, screening decisions, extracted study summary statistics, Risk-of-Bias appraisals, GRADE evidence tables, meta-analyses, and manuscript drafts), alongside minimal user account records (email, display name,

salted adaptive password hash, AES-256-GCM encrypted TOTP seed, session/audit metadata, and Stripe billing identifiers). **CoreSR never connects to hospital EMRs and strictly prohibits Protected Health Information (PHI) or identifiable patient records** under PHIPA, HIPAA, or GDPR Article 9 (Terms of Service §6.1; DPA §2.3).

Q4. Where is Customer Data hosted, backed up, and processed?
Answer: Primary production compute, relational databases, and AES-256 encrypted file volumes are hosted in **Helsinki, Finland (EU / EEA)** at **Hetzner Online GmbH** (ISO/IEC 27001 certified data centres). Primary GPU inference (`TITAN-SR` , `GLASS` , `ROBERT`) and verified off-site encrypted backups operate over an authenticated point-to-point encrypted tunnel on CoreSR-operated infrastructure in **Ontario, Canada** (covered by PIPEDA and EU Commission Adequacy Decision 2002/2/EC, reaffirmed 2024), with asymmetrically pre-encrypted backup ciphertext replicated to Google Cloud/Workspace and Compliance-mode WORM Object Lock storage on Backblaze B2 (neither cloud storage provider holds decryption keys). For **Team** and **Institution** customers requiring US-only residency, a dedicated **US-region deployment option** is available under an Institutional Order Form.

Q5. Do you offer a PIPEDA / GDPR Data Processing Addendum (DPA) and advance Subprocessor notice?
Answer: Yes. CoreSR publishes a pre-drafted, ready-to-countersign **PIPEDA & GDPR Article 28 Data Processing Addendum** (incorporating the 2021 EU Standard Contractual Clauses, Module Two, and UK IDTA) as a signed letterhead PDF at https://coresr.ai/api/public/dpa, alongside our **Authorized Subprocessor Register** at https://coresr.ai/api/public/subprocessors. CoreSR contractually commits to **at least 30 calendar days' advance written notice** before adding or materially changing any Subprocessor, with a 14-day objection and pro-rata refund termination right (DPA §5.3).

Q6. How do data export, retention, and deletion work when a user or institution leaves?
Answer: Users and workspace owners have immediate self-service **One-Click Full Account ZIP Export** (`Account Settings → Full Data Export`), exporting all references, screening decisions with pinned model versions, extractions, Risk-of-Bias assessments, GRADE tables, manuscripts, and lossless JSON snapshots. Initiating **Self-Service Account Deletion** enters a **7-day revocable grace window** (revoking other sessions immediately), after which automated purge workflows permanently delete live database records and uploaded files, retaining only an irreversible SHA-256 email digest in the deletion audit log. Pre-encrypted off-site backup archives automatically expire and age out within their rolling retention window (30–35 days).

Section C — Encryption, Network Security & Administrative Access

Q7. How is Customer Data encrypted in transit and at rest?
Answer:

- **In transit:** Public HTTPS enforces **TLS 1.3** (modern AEAD cipher suites only) with **HSTS preload** and **mutual TLS (mTLS)** origin authentication; private east–west traffic between European Union and Canadian infrastructure traverses an authenticated encrypted tunnel.
- **At rest:** Production databases, uploaded study PDFs, and storage volumes reside on **AES-256** encrypted volumes with hardware-isolated keys separated from the boot disk. Passwords use salted adaptive cryptographic hashing; MFA seeds use AES-256-GCM; session tokens are stored solely as SHA-256 digests; off-site backups are asymmetrically encrypted prior to cloud replication.

Q8. How is administrative access to production systems governed, and what MFA controls exist for customer workspaces?
Answer:

- **Infrastructure & Console Access:** All cloud, source-control, and production administrative access is restricted to individually attributed accounts under strict least-privilege RBAC, 100% hardware/TOTP multi-factor authentication, public-key-only server access (password sign-in disabled), zero shared credentials, and quarterly access reviews (Privileged Access Register inspectable under NDA).
- **Customer Workspace Controls:** Team and Institution Workspace Owners and Admins can enforce **mandatory TOTP 2FA** across all workspace members, shorten **maximum session lifetimes** (down to 8 hours), inspect member 2FA enrollment

status, and **revoke active member sessions** immediately.

Section D — AI Governance, Model Training & Third-Party Egress (Controls 4.1–4.6)

Q9. Is our organization's content (protocols, unpublished guidelines, PDFs, screening decisions, extractions, or drafts) ever used to train or fine-tune CoreSR's AI models?
Answer: No. For all **Team** and **Institution** workspaces, Customer Content is **never** used to train, fine-tune, distil, or evaluate CoreSR models (enforced in server export pipelines and warranted in DPA §3.1 and Terms §8.1). The sole exception is where an institutional research partner elects to execute a separate, bilaterally signed written Research / Co-Development Agreement scoped strictly to its own designated dataset — never under standard subscriptions. Standalone Individual/Free users also have a one-click **"Exclude my reviews from model training"** toggle in Account Settings.

Q10. Does CoreSR send prompts, uploaded PDFs, or review data to OpenAI, Anthropic, Google Gemini, DeepSeek, or other third-party commercial LLM vendors?
Answer: No, with one exception. Third-party commercial LLM egress is blocked in server code for all Team and Institution workspaces. The exception, effective November 17, 2026, is Kaplan–Meier figure reading: when a reviewer runs KM Cor on a 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; no prompt, uploaded PDF or other review data is sent, and a figure can instead be traced by hand without AI. For everything else, CoreSR runs its own domain-specific models — **TITAN-SR** (`v6-2026-10-15`), **GLASS v3** (6 task specialists, 65,536-token context), and **ROB BERT v6.1** — on CoreSR-operated infrastructure (EU & Canada) and, only when CoreSR's private servers are down or unavailable, stateless zero-retention GPU containers executing CoreSR's own model weights (RunPod, Inc., SOC 2 Type II). Server policy blocks every other call to external commercial LLM APIs before any outbound connection can open. CoreSR never uses DeepSeek or any China-domiciled AI provider on any tier.

Q11. What were TITAN-SR and GLASS trained on, and what is your licensing basis?
Answer: TITAN-SR v6 was trained on 3,097,955 publicly distributed scholarly bibliographic records (titles, abstracts, and indexing metadata from PubMed/MEDLINE, Europe PMC, OpenAlex CC0, Crossref, and ClinicalTrials.gov) paired with published systematic-review inclusion outcomes across 19,926 review groups on Apache-2.0/MIT encoder backbones. **GLASS v3** is built on permissively licensed (**Apache-2.0**) open-weight foundation weights, domain-adapted on the 3,097,955 screening corpus and fine-tuned on open-access biomedical full texts (PMC Open Access / CC-BY / CC0), published methodological guidance, and internally verified structured appraisal worksheets — never patient PHI or pirated publisher archives.

Q12. Are prompts or study excerpts stored in AI inference logs? How is model provenance tracked for peer review?
Answer: Operational telemetry records only operation counts, model identifiers, token totals, latency, and estimated compute cost — **never prompts, responses, URLs, or study text** — and stateless GPU containers retain zero bytes post-response. Every AI output pins its exact model release tag (`v6-2026-10-15` , `dual3-2026-07-12` , `glass-v3` , `v6.1`) in project records and CSV/ZIP exports, backed by our public Model Version Changelog (https://coresr.ai/api/public/model-changelog).

Q13. How does CoreSR prevent AI hallucinations from corrupting a guideline or systematic review?
Answer: Human-in-the-loop verification is enforced architecturally: AI screening classifications are stored separately from human adjudications; GLASS extractions require a locatable verbatim PDF quote and page number (unlocatable quotes are rejected automatically) and remain unverified until human reviewer sign-off; RoB 2 domain algorithms and quantitative CoreGRADE domains (**Risk of Bias**, **Imprecision**, **Publication Bias**) are computed **100% deterministically in code**, never estimated by generative prose.

Section E — Resilience, Continuity, Incident Response & Accessibility

Q14. What are your backup architecture, RPO/RTO, and operational continuity safeguards?
Answer: CoreSR maintains a **3-2-1-1 backup architecture** (**RPO ≤ 24 hours · RTO ≤ 24 hours**): daily database snapshots and file archives are replicated off-host, asymmetrically encrypted before cloud upload, verified continuously via automated decryption and catalog integrity checks, and stored across Google Cloud/Workspace and Compliance-mode WORM Object Lock on Backblaze B2 (write-only credentials incapable of deletion), validated by quarterly clean-room restore drills. Security patching, certificate renewal, health watchdogs, and alerting operate autonomously, supported by a formally designated executive Continuity Custodian with escrowed emergency recovery procedures and an optional **Source-Code & Container Escrow Rider** for Institution agreements.

Q15. What is your security incident response and breach notification SLA?
Answer: CoreSR maintains a formal Incident Response Plan with scenario playbooks (Account Compromise, Data Exposure, Key Compromise, Ransomware) and real-time automated security anomaly alerting. Confirmed Personal Data Breaches are notified to affected customers **without undue delay and within 72 hours** (targeting **24 hours** for Institution customers) per GDPR Art. 33 and PIPEDA s. 10.1 (security@coresr.ai).

Q16. Does CoreSR conform to WCAG 2.1 Level AA / AODA / Section 508 accessibility standards?
Answer: Yes. CoreSR conforms to **WCAG 2.1 Level A and Level AA** (zero automated accessibility violations across 20 audited public and core workflow views, full keyboard operability, ≥ 4.5:1 text contrast, ≥ 3:1 focus indicators, unique route titles, and accessible HTML/CSV data tables paired with every statistical figure). Our complete **ITI VPAT® 2.5 Accessibility Conformance Report (PDF)** is published at https://coresr.ai/api/public/vpat.

ISSUED ON BEHALF OF CORESR INC.

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