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

PUB-005 · Version 1.3
Effective: 2026-11-17
Public Procurement & Governance

Security Controls Statement

ISMS Framework: Internal programme modelled on ISO/IEC 27001:2022
Security, Privacy & Legal Contacts: security@coresr.ai · legal@coresr.ai

This Security Controls Statement sets out the technical and organisational controls implemented in production across **CoreSR** (https://coresr.ai). The four control areas below match Section 4 of the public Security & Trust page (https://coresr.ai/security) word-for-word and align with Annex II of the CoreSR PIPEDA & GDPR Article 28 Data Processing Addendum (https://coresr.ai/api/public/dpa).

1. Encryption

- **In transit:** All public HTTPS traffic is terminated at the edge under TLS 1.3 (minimum TLS 1.2 with modern AEAD cipher suites only), HSTS preload, and mutual TLS (mTLS) origin authentication between the CDN edge and CoreSR origin servers. Private east–west traffic between European Union primary production and Canadian GPU and disaster-recovery infrastructure traverses an authenticated point-to-point encrypted tunnel.
- **At rest on production storage:** Production relational databases, uploaded study PDFs, extracted review work product, and staging archives reside on AES-256 encrypted storage volumes in ISO/IEC 27001 certified European Union data centres, with cryptographic keys isolated from the operating-system boot volume.
- **Application credentials & secrets:** Account passwords are stored exclusively as salted adaptive cryptographic hashes; TOTP multi-factor authentication seeds are encrypted at rest with AES-256-GCM using keys separated from the database; browser session tokens are persisted server-side solely as irreversible SHA-256 digests; optional customer bibliographic API keys are encrypted at rest and masked across all API responses.
- **Off-site backups:** Every database snapshot and file archive is asymmetrically encrypted prior to off-site replication using modern authenticated cryptography. Cloud storage subprocessors receive ciphertext only and never hold decryption keys.

2. Network Security

- **Origin isolation & edge protection:** Public ingress routes exclusively through global Anycast WAF and DDoS mitigation. Host-level ingress policies default-deny inbound connections, permit HTTPS solely from authenticated edge origins verified by mutual TLS client certificates, and drop direct-to-origin probes.
- **Internal tier isolation:** Relational database and internal application workers bind strictly to non-public internal interfaces and are never exposed to the public internet.
- **Hardened administrative access:** Infrastructure administration requires individual attributed accounts, cryptographic public-key authentication (password sign-in disabled), brute-force intrusion blocking, and 100% multi-factor authentication across all cloud, source-control, and administrative consoles, reviewed quarterly.

- **Outbound SSRF & HTTP header defences:** Server-side URL fetchers validate protocols and resolved network ranges to block internal, private, link-local, and cloud-metadata destinations. Application responses enforce strict Content-Security-Policy, X-Frame-Options, X-Content-Type-Options, Referrer-Policy, and Permissions-Policy headers.

3. Vulnerability Management

- **Continuous CI/CD security gates:** Every pull request, main-branch release, and scheduled weekly audit runs blocking static application security testing (SAST), software composition analysis (SCA) across application and ML dependencies, and automated repository secret detection.
- **Automated operating-system patching:** Production hosts run enterprise long-term-support Linux distributions with extended security maintenance, automated daily security errata patching, and immediate operator alerting on update anomalies.
- **Recurring dynamic scanning & vulnerability disclosure:** Automated dynamic application security testing (DAST) scans run weekly against production, and external security researchers may report findings under safe-harbor terms via https://coresr.ai/.well-known/security.txt (security@coresr.ai).
- **Binding remediation SLAs & penetration testing:** Validated vulnerabilities are remediated under written SLAs: Critical ≤ 7 calendar days, High ≤ 30 calendar days, and Medium ≤ 90 calendar days. Independent third-party penetration testing is commissioned upon institutional procurement requirement or at SOC 2 Type II kickoff (scheduled Q1 2027), whichever occurs first.

4. Data Loss Prevention

- **Published-literature boundary (No PHI):** CoreSR is engineered strictly for published scientific literature and systematic-review work product; uploading identifiable patient health records or Protected Health Information (PHIPA / HIPAA / GDPR Art. 9) is contractually prohibited and excluded by design.
- **Third-party commercial LLM egress blocked (Team & Institution):** For all Team and Institution workspaces, server-side policy blocks outbound calls to external commercial LLM APIs in code and permanently excludes workspace content from model training, confining inference to CoreSR-operated infrastructure (and, only when those private servers are down, stateless zero-retention GPU containers running CoreSR's own models). One exception: when a reviewer runs KM Cor on a Kaplan–Meier figure, the cropped figure, its caption and the review's subgroup names are sent to Google's Gemini vision model.
- **Tenant isolation, MFA & session governance:** Every non-public API route enforces authenticated session validation, per-project and workspace role-based access checks backed by automated cross-tenant isolation tests, CSRF protection, and adaptive rate limiting. Team and Institution administrators can mandate workspace-wide TOTP 2FA, enforce shortened session lifetimes, and revoke active member sessions immediately.
- **Audit trail, portability, erasure & immutable resilience:** Security-sensitive events write to an append-only audit trail with real-time anomaly alerting. Customers have self-service One-Click Full Account ZIP Export, revocable account deletion followed by automated live purge, and geo-redundant encrypted backups including Compliance-mode WORM Object Lock immutability (RPO ≤ 24 hours, RTO ≤ 24 hours).

Companion Procurement & Assurance Instruments

Instrument	Public URL (PDF)
Institutional Vendor Information Pack	https://coresr.ai/api/public/vendor-pack
PIPEDA & GDPR Article 28 Data Processing Addendum	https://coresr.ai/api/public/dpa

Instrument	Public URL (PDF)
Authorized Subprocessor Register	https://coresr.ai/api/public/subprocessors
AI Governance & Provenance Specification (Controls 4.1–4.6)	https://coresr.ai/api/public/ai-governance
VPAT® 2.5 WCAG 2.1 Level A & AA Conformance Report	https://coresr.ai/api/public/vpat

Detailed internal ISMS policy documents, Statement of Applicability schedules, and technical diligence briefings are provided to institutional security teams under mutual NDA via security@coresr.ai.

ISSUED ON BEHALF OF CORESR INC.

Tyler Pitre, MD PhD FRCPC

Founder and Chief Executive Officer, CoreSR Inc.

tyler.pitre@coresr.ai · coresr.ai