Srasta — CAIQ Lite Questionnaire Responses

Version: 1.0 **Last Updated:** 2026-04-29 **Framework:** Cloud Security Alliance Cloud Controls Matrix (CCM) v4 — Lite subset (~45 questions across 17 domains) **Service:** Srasta Platform v1.x

This document provides pre-drafted responses to the CAIQ Lite question set so prospective customers can evaluate Srasta without waiting for a fresh round-trip on every security questionnaire. For the full CAIQ v4 (~261 questions), additional responses can be provided on request — most reduce to the same architectural invariants captured below.

How to read this document

- **Response** is one of: ✅ Yes / ⊘ No / 🔄 Partial / ⊘ N/A.
- **Implementation** is a one-line statement of how the control is satisfied.
- **Evidence** points at code, doc, or runtime artifact (every evidence pointer is reproducible from the running system).
- **Notes** flag anything that's edition-specific, customer-side, or roadmap.

For the SOC 2 Common Criteria mapping behind these responses, see `docs/security/controls-matrix.md`. For the architectural posture referenced throughout, see `docs/security/architecture.md`.

Architectural anchor

Srasta is **self-hosted**: every Srasta service runs inside the customer's infrastructure perimeter. Customer end-user data (documents, prompts, responses, audit records) **never leaves that perimeter** in normal operation. Many CAIQ controls that are heavyweight for traditional SaaS reduce to **"customer- operated under the shared-responsibility model"** in Srasta's case — that is reflected in the responses below.

The Gandiva-side license-server (the only Gandiva-operated runtime surface) is in scope only for license issuance + the narrow personal-data set described in `docs/legal/privacy-policy.md` Section 3.

Application & Interface Security (AIS)

#	Question	Response	Implementation	Evidence
AIS-01	Are application security policies established and maintained?	✅ Yes	Operator-configurable policy profiles (`baseline` / `hipaa` / `pci` / `fedramp`) enforce input/output policy at the rag-api + tool-gateway layers. Baseline includes PII detection + redaction; stricter profiles add PHI/PAN blocking.	the policy engine config; `setup/tool_policy.yaml`
AIS-02	Are application security requirements identified and documented?	✅ Yes	Stated in `docs/security/controls-matrix.md` (CC5, CC6, CC7) and `docs/security/architecture.md` .	`docs/security/controls-matrix.md`
AIS-03	Are inputs validated at all entry points?	✅ Yes	srasta-api gateway is the single external entry point. All non-/health, non-/oauth requests pass through auth + authz + license + rate-limit middleware before reaching upstream services. Internal services reject any request lacking a valid HMAC-signed `X-Srasta-Signature` .	the gateway request middleware; the HMAC signer for forwarded principal headers
AIS-04	Is secure software development lifecycle (SDLC) maintained?	✅ Yes	Git-based SDLC with required GitLab MR review, CI gates (ruff, validate-version, validate-no-hardcoded-ips, validate-fresh-install, grype, syft, cosign), digest-pinned dependencies, signed releases.	`.gitlab-ci.yml` ; `release/services.yaml` ; the no-tactical-drift policy

Audit & Assurance (A&A)

#	Question	Response	Implementation	Evidence
A&A-01	Is an internal audit program in place?	🔄 Partial	Engineering audits run continuously (CI gates + admin UI security panel + audit-feed live tail + chain-integrity verification). Formal annual internal audit lands during SOC 2 Type I prep.	the admin security panel; the audit-log chain-verification routine
A&A-02	Are independent audits / certifications maintained?	📅 Roadmap	SOC 2 Type I targeted post-seed-close (see `docs/strategy/seed-momentum-plan.md` ticket #7). Vendor will be Drata. No certification yet.	`docs/strategy/seed-momentum-plan.md`
A&A-03	Can customers audit the platform on-premises?	✅ Yes	The platform runs inside the customer's perimeter; the customer can audit it directly. Source for the customer's deployment is image-pinned + cosign-signed. SBOMs available per release.	`scripts/verify-release-bundle.sh` ; per-release CI artifacts

Business Continuity Management & Operational Resilience (BCR)

#	Question	Response	Implementation	Evidence
BCR-01	Is a Business Continuity Plan (BCP) maintained?	🔄 Partial	Runtime backup/restore/rollback tooling shipped (`scripts/backup.sh` , `restore.sh` , `rollback.sh`). Formal BCP doc with RTO/RPO targets lands with `docs/security/bcp-dr-plan.md` (planned alongside SOC 2 Type I prep).	`scripts/backup.sh` ; `scripts/rollback.sh`
BCR-02	Are backups encrypted?	🔄 Partial	Backups land in MinIO inside the customer's perimeter. At-rest encryption depends on the underlying storage (host disk encryption, MinIO server-side encryption, or operator-configured external S3 with SSE). Customer-controlled.	`scripts/backup.sh` ; customer infra config
BCR-03	Are backups tested?	✅ Yes	`scripts/verify-backup.sh` performs a full restore into an isolated container. Recommended quarterly; CI runs a smoke version on every release.	`scripts/verify-backup.sh`
BCR-04	What is the RTO / RPO target?	🔄 Partial	License-server side (Gandiva-operated): RTO 4 hours, RPO 1 hour. Customer-deployed Srasta side: customer determines based on their backup cadence. Default backup cadence is on-upgrade with 5-most-recent retention.	`docs/security/bcp-dr-plan.md` (planned)

Change Control & Configuration Management (CCC)

#	Question	Response	Implementation	Evidence
CCC-01	Are changes managed through a defined process?	✅ Yes	All changes flow through GitLab MR with required review. Every commit references a GitLab issue. CI pipeline gates merges; releases require signed tags.	the commit-issue-ref policy; `.gitlab-ci.yml` ; `git log`
CCC-02	Is configuration baselined?	✅ Yes	Customer compose generated from `release/services.yaml` (single source of truth). Image digests pinned in customer compose. Operator config persisted in `platform_config` (Postgres) — `Mandate 2: operator-configurable, never code-baked` .	`release/services.yaml` ; the customer-compose generator
CCC-03	Are changes tested before production?	✅ Yes	CI gates: ruff, validate-version, validate-no-hardcoded-ips, validate-fresh-install (every dev pipeline does a clean-install smoke), grype CVE scan, syft SBOM, cosign sign + verify. gdlab.net is the canary deployment; releases land there before any other customer.	`.gitlab-ci.yml` ; the audit-thesis design principle

Cryptography, Encryption & Key Management (CEK)

#	Question	Response	Implementation	Evidence
CEK-01	Is data encrypted in transit?	✅ Yes	TLS 1.2+ for all external endpoints (nginx via Let's Encrypt wildcard or operator BYOC). Service-to-service uses HMAC-signed forwarded headers + Docker bridge isolation.	`setup/nginx/` ; the gateway authentication layer
CEK-02	Is data encrypted at rest?	🔄 Partial	Customer-deployed Srasta: at-rest encryption depends on host disk encryption (LUKS / FileVault / hardware FDE) — customer-controlled. Gandiva-side license-server data: AES-256 (cloud provider managed key).	`docs/security/architecture.md` Section "Where customer data lives"
CEK-03	Are cryptographic keys rotated?	✅ Yes	`scripts/rotate-secrets.sh` rotates Tool Gateway API keys, Langfuse secrets, MinIO credentials with .env backup + service restart + post-rotation smoke. License RSA keypair rotation is a Gandiva-ops procedure (planned annually).	`scripts/rotate-secrets.sh` ; `docs/compliance/key-rotation.md`
CEK-04	What signature algorithms are used?	✅ RS256 / SHA-256	License JWTs: RS256 (RSA 2048). Audit log hash chain: SHA-256. Forwarded headers: HMAC-SHA256. Container images: Cosign keyless via Sigstore.	the license-validation library; the shared audit-log module; the gateway authentication layer

Datacenter Security (DCS)

#	Question	Response	Implementation	Evidence
DCS-01	Is physical access to datacenters controlled?	⊘ Customer	Srasta runs inside the customer's infrastructure; physical security is customer-controlled. Gandiva-side license-server runs in a major cloud provider with their physical security posture.	Customer-attested

Data Security & Privacy Lifecycle Management (DSP)

#	Question	Response	Implementation	Evidence
DSP-01	Is data classified?	✅ Yes	PII / PHI / PAN / CVV are detected by the policy engine on every input + output (via regex patterns in the PII/PHI/PAN detection library). Classification feeds redaction / blocking decisions per the active policy profile.	the PII/PHI/PAN detection library; the policy engine
DSP-02	Is data handling appropriate to classification?	✅ Yes	PHI under HIPAA profile is redacted before forwarding. PAN/CVV under PCI profile is hard-blocked. Audit entries record classification outcome (`phi_detected` , `pan_detected`).	`audit/policy.py:PolicyConfig.validate_request`
DSP-03	Is data minimized?	✅ Yes	Prompts and responses are not persisted by default. Held in memory only for the duration of the request. Optional Langfuse trace storage opt-in via `--profile langfuse` . Operational telemetry (counts + shapes only) is the only default outbound channel; payload schema enumerated in `docs/legal/privacy-policy.md` Section 3.4 — no prompt / response / document / audit-log content ever leaves the customer perimeter.	`docs/security/architecture.md` "Where customer data lives"; `docs/legal/privacy-policy.md` Section 3.4
DSP-04	Are data subjects' rights honored?	✅ Yes	Customer is the data controller for end-user data; rights handling lives in their privacy policy. For Gandiva-controlled data (license, support, marketing), see `docs/legal/privacy-policy.md` Section 8.	`docs/legal/privacy-policy.md`
DSP-05	Are AI training uses disclosed?	✅ Yes (negative)	Customer end-user data is **never** used to train Gandiva or third-party models. All inference runs in customer infrastructure (vLLM / Ollama / TEI) or routes to provider APIs the customer chose, where the provider's TOS governs training reuse.	`docs/security/architecture.md` "What Srasta does NOT do"
DSP-06	Is data deletion supported?	✅ Yes	Customer side: standard infra-level deletion (volume removal, Postgres truncate). Audit log	the audit-log rotation routine

#	Question	Response	Implementation	Evidence
			rotation on `AUDIT_LOG_RETENTION_DAYS` boundary. Gandiva side: per `docs/legal/privacy-policy.md` Section 3, retention is bounded by purpose.	

Governance, Risk Management & Compliance (GRC)

#	Question	Response	Implementation	Evidence
GRC-01	Is an information security program established?	🔄 Partial	Operational controls in place (covered by this matrix). Formal Information Security Policy doc lands with SOC 2 Type I prep.	`docs/security/controls-matrix.md`
GRC-02	Are compliance frameworks tracked?	✅ Yes	SOC 2 CC matrix (`docs/security/controls-matrix.md`). HIPAA / PCI / FedRAMP multi-framework reference at `docs/compliance/controls-matrix.md` (post-seed scope per Mandate 5).	`docs/security/controls-matrix.md`
GRC-03	Is risk assessment performed?	🔄 Partial	Continuous tactical risk surfacing via security panel + audit feed. Formal risk register lands with SOC 2 Type I prep (Q3 target).	the admin security panel; `docs/security/controls-matrix.md` row CC3.2

Human Resources Security (HRS)

#	Question	Response	Implementation	Evidence
HRS-01	Are background checks performed for employees with production access?	📅 Post-Type-I	Solo-founder pre-seed. Background-check program tied to first hire policy.	—
HRS-02	Is security awareness training delivered?	📅 Post-Type-I	Same as above; tied to first hire.	—

Identity & Access Management (IAM)

#	Question	Response	Implementation	Evidence
IAM-01	Is multi-factor authentication (MFA) supported?	✅ Yes	MFA is enforced at the Customer's OIDC provider (Zitadel / Okta / Azure AD / Google). Srasta delegates authentication entirely to OIDC; the platform validates the resulting JWT.	the gateway OIDC auth provider
IAM-02	Is role-based access control (RBAC) implemented?	✅ Yes	Roles assigned via the OIDC provider arrive as JWT claims and are mapped to Srasta capabilities. Per-role model whitelist is enforced at the gateway (`platform_config.model_access`).	the gateway authorization layer; the gateway per-role model authorization
IAM-03	Is privileged access logged?	✅ Yes	Every admin action emits an audit event via `_append_audit` with hash-chain integrity. Live-tail in admin UI Activity tab; canonical taxonomy filters to `event_class=admin` .	the admin audit emission helper; #147 taxonomy
IAM-04	Is session timeout enforced?	✅ Yes	RAG-API session state purged after `SESSION_TTL_HOURS` (default 24h) idle. Token validity controlled by OIDC provider.	the RAG service; OIDC provider config
IAM-05	Are service / API credentials managed?	✅ Yes	Tool-gateway API keys in `TOOL_GATEWAY_API_KEYS` (rotatable via `scripts/rotate-secrets.sh`). License JWT in `platform_config` (admin UI editable).	`scripts/rotate-secrets.sh` ; `docs/roadmap/licensing-roadmap.md`

Interoperability & Portability (IPY)

#	Question	Response	Implementation	Evidence
IPY-01	Is data exportable in standard formats?	✅ Yes	Audit log: JSONL. Backup bundle: pg_dump SQL + Milvus metadata JSON + .env. CSV export from admin UI Activity tab (via the audit feed).	the shared audit-log module; `scripts/backup.sh` ; admin UI Activity tab
IPY-02	Is the platform vendor-portable?	✅ Yes	Docker compose deployment is portable across single-host installs. Helm chart supports K8s deployments. No managed-SaaS lock-in.	`chart/` ; `docker-compose.yml`

Infrastructure & Virtualization Security (IVS)

#	Question	Response	Implementation	Evidence
IVS-01	Is network segmentation enforced?	✅ Yes	All Srasta services run on a private Docker bridge network (`srasta`). Only ports listed in `docker-compose.yml` are exposed. K8s deployment supports NetworkPolicy.	`docker-compose.yml:networks.srasta` ; `chart/templates/srasta-admin/networkpolicy.yaml`
IVS-02	Are workloads isolated?	✅ Yes	Each service runs in its own container. All four audit-writing services run as the same non-root UID (1000) for shared-volume access without a UID race (#164). K8s deployments add `runAsUser: 1000` , `runAsGroup: 1000` , `fsGroup: 1000` .	`chart/templates/*/deployment.yaml` ; `docker-compose.yml:srasta-audit-init`
IVS-03	Is least-privilege enforced for service accounts?	✅ Yes	K8s `automountServiceAccountToken: false` on all pods; namespace-scoped RBAC (no ClusterRole).	`chart/templates/*/deployment.yaml`

Logging & Monitoring (LOG)

#	Question	Response	Implementation	Evidence
LOG-01	Is audit logging implemented?	✅ Yes	Hash-chained JSONL audit log covers every API request, auth event, authz denial, rate-limit, license denial, admin mutation. Canonical event taxonomy (#147) classifies events into auth/inference/tool/admin.	the shared audit-log module; the audit-event canonical classifier
LOG-02	Is audit log integrity protected?	✅ Yes	SHA-256 chain (`prev_hash + entry_json`). `verify_chain()` returns deterministic errors on any tamper. Files mode 0644 owned by `srasta:1000` (#164).	the audit-log chain-verification routine
LOG-03	Is audit log retention configurable?	✅ Yes	`AUDIT_LOG_RETENTION_DAYS` env var (default 90). Atomic rotation rewrite preserves chain continuity.	the audit-log rotation routine; `.env.example`
LOG-04	Is audit log forwardable to a SIEM?	✅ Yes	Optional `--profile audit` enables Vector forwarder. Supported sinks: S3 / Splunk HEC / RFC 5424 syslog / generic HTTP webhook. HMAC-signed for integrity.	`setup/vector/vector.toml` ; the audit-forwarder daemon
LOG-05	What events are logged?	✅ Comprehensive	All API requests; auth events (success/fail); authz decisions; rate-limit hits; license events (granted/renewed/cleared/denied); policy enforcement decisions (PHI/PAN); admin mutations; backup + restore operations; ingest runs.	the shared audit-log module; the admin audit emission helper calls
LOG-06	Is operational telemetry sent to the vendor?	🔄 Partial	Daily phone-home from each Srasta install (default ON for both community + enterprise editions, opt-out via `SRASTA_TELEMETRY=off`) ships counts + shapes only — never customer data. Schema in `docs/legal/privacy-policy.md` Section 3.4. Phase 2 of #173 wires the receiver.	`docs/legal/privacy-policy.md` Section 3.4; `docs/roadmap/self-serve-trial-roadmap.md` D5

Security Incident Management (SEF)

#	Question	Response	Implementation	Evidence
SEF-01	Is an incident response plan documented?	🔄 Partial	Multi-framework playbooks exist at `docs/compliance/incident-response.md` (P1–P3 severity, detect/contain/investigate/recover/post-mortem). Seed-stage runbook + escalation matrix lands during SOC 2 Type I prep.	`docs/compliance/incident-response.md`
SEF-02	Is breach notification timeline defined?	✅ Yes	72-hour notification per GDPR Art. 33 for Gandiva-controlled-data breaches. Customer-deployment breaches: customer is responsible for detection and notification; Gandiva provides technical assistance on request.	`docs/legal/privacy-policy.md` Section 7; `docs/compliance/incident-response.md`
SEF-03	How are vulnerability reports handled?	✅ Yes	Reports to `security@gandiva.tech`. Acknowledgement within 2 business days; fix timeline within 5 business days for critical issues.	`docs/security/responsible-disclosure.md` (planned); `docs/compliance/incident-response.md`

Supply Chain Management (STA)

#	Question	Response	Implementation	Evidence
STA-01	Is a sub-processor list maintained?	✅ Yes	`docs/legal/privacy-policy.md` Section 6 + DPA Annex 2 enumerate sub-processors. Customers notified before material changes per DPA notification clause.	`docs/legal/privacy-policy.md` Section 6
STA-02	Is supply-chain integrity verified?	✅ Yes	Cosign-keyless image signing via Sigstore (GitLab OIDC). Per-image SBOMs (Syft). Customer-side `scripts/verify-release-bundle.sh` validates digests + signatures before extraction.	`release/services.yaml`; `scripts/verify-release-bundle.sh`; `.gitlab-ci.yml`
STA-03	Are open-source dependencies tracked?	✅ Yes	All deps pinned to specific versions in `pyproject.toml` files; image digests pinned in customer compose; Syft SBOMs published per release. Renovate / dependabot review monthly.	Per-service `pyproject.toml`; CI SBOM artifacts

Threat & Vulnerability Management (TVM)

#	Question	Response	Implementation	Evidence
TVM-01	Is vulnerability scanning performed?	✅ Yes	Grype scans every container image in CI; blocks merge on CRITICAL findings. Customers can re-scan via `cosign download sbom srasta/<image>:<version>`.	`.gitlab-ci.yml` grype job; per-release SBOM
TVM-02	Is patch management documented?	✅ Yes	Critical (CVSS ≥ 9.0) patches: 48 hours target. High-severity: 30-day target. Operators apply via `scripts/upgrade.sh` with rollback safety net.	`docs/compliance/sla.md` Section 4; `scripts/upgrade.sh`; `scripts/rollback.sh`
TVM-03	Are penetration tests performed?	📅 Post-Type-I	Annual third-party pentest is committed for Q3 2026 alongside SOC 2 Type I prep.	`docs/strategy/seed-momentum-plan.md`

Universal Endpoint Management (UEM)

#	Question	Response	Implementation	Evidence
UEM-01	Are endpoints managed?	⊘ Customer	Srasta endpoints (containers / pods) run inside the customer's environment. Endpoint management (host hardening, EDR, patching) is customer-operated under the shared-responsibility model.	`docs/compliance/shared-responsibility.md`

Notes on response semantics

- 🔄 **Partial** is honest: it means the control is satisfied for the running system but the **formal documentation / process artifact** is still in flight (typically tied to SOC 2 Type I prep or the first hire).
- 📅 **Post-Type-I** is similarly honest: the control is genuinely not yet implemented and is timeline-bound to seed-close.
- ⊘ **Customer** is not a deflection; it reflects the architectural reality that Srasta runs in the customer's perimeter. The shared-responsibility document (`docs/compliance/shared-responsibility.md`) makes this allocation explicit.

Document index

Topic	Document
Architectural data-flow + trust boundaries	`docs/security/architecture.md`
SOC 2 Common Criteria mapping	`docs/security/controls-matrix.md`
Privacy policy + Gandiva-controlled data	`docs/legal/privacy-policy.md`
Multi-framework controls (HIPAA / PCI / FedRAMP)	`docs/compliance/controls-matrix.md`
Shared responsibility model	`docs/compliance/shared-responsibility.md`
Incident response playbooks	`docs/compliance/incident-response.md`
DPA template	`docs/compliance/dpa.md`
BAA template (HIPAA)	`docs/compliance/baa-template.md`
Audit logging detail	`docs/compliance/audit-logging.md`
Backup + recovery detail	`docs/compliance/backup-recovery.md`
Key rotation procedures	`docs/compliance/key-rotation.md`

This document is a snapshot. As controls evolve, this matrix is revised together with `docs/security/controls-matrix.md`. The authoritative status of any single control is whichever document was updated most recently — both are auto-tracked via git.