Srasta — Privacy Policy

Effective Date: 2026-06-22 **Last Updated:** 2026-06-22 **Version:** 1.1 (template — pending legal review)

> **LEGAL NOTICE:** This is a template policy adapted from standard self-hosted-software privacy postures. It has not yet been reviewed by outside counsel. Customers and prospects may rely on the architectural invariants stated here (the data-flow boundary), but should treat the legal-language formulations as draft. Outside-counsel review is committed by Type-I prep (see `docs/security/soc2-roadmap.md`, planned).

---

1. Who We Are

Gandiva Tech, Inc. ("Gandiva", "we", "us") publishes the **Srasta Platform** ("Srasta") — a governed AI execution platform that runs inside customer-controlled infrastructure.

- **Marketing site:** https://srasta.ai
- **Privacy contact:** privacy@srasta.ai
- **Security contact:** security@srasta.ai

This Privacy Policy describes how Gandiva handles personal data in its capacity as the publisher of Srasta, the operator of the Gandiva-side license-server, and the operator of marketing and support channels. It does **not** describe how customers handle the personal data they process inside their own Srasta deployment; that is the customer's responsibility under their own privacy policy and the DPA they execute with us.

---

2. The Architectural Invariant

Srasta is **self-hosted**: every Srasta service runs inside the customer's infrastructure perimeter — their VPC, their cluster, or their on-prem hosts. The customer's **end-user data** (documents, prompts, responses, audit records, identity records) **never leaves that perimeter** in normal operation.

In particular:

- Gandiva does not receive, route, or store the customer's end-user data. Customer data flows are documented in `docs/security/architecture.md`.
- Gandiva does not have backdoor access into customer deployments.
- The Srasta runtime transmits **operational metadata only** (install events, version, cluster size, feature-touch counts, error class counts) to Gandiva by default — see Section 3.4 for the full schema, lawful basis, and opt-out mechanism. The telemetry payload **never contains** prompt content, model response content, document content, audit-log content, end-user identifiers, or any other customer-controlled data.
- All inference (LLM calls) is routed by the customer's own configuration to either local models running in customer infrastructure or to third-party providers the customer chooses (e.g. OpenAI, Anthropic) — Gandiva does not sit in this path.

For the data-flow diagram, see `docs/security/architecture.md`. Trust-boundary table, outbound-channel inventory, and per-data-class storage map are the authoritative artifacts; this policy text references them.

---

3. The Personal Data Gandiva Does Process

Despite the architectural invariant above, there are narrow surfaces where Gandiva itself receives personal data — limited to operating the Gandiva-side license server, support channels, and the marketing site. Each is enumerated here.

3.1 License-server data

Purpose: to issue, renew, and revoke license keys for paying customers.

Field	Source	Lawful basis (GDPR)	Retention
Customer organization name	Customer order form	Contract performance (Art. 6(1)(b))	Duration of customer relationship + 7 years (tax/accounting)
Customer billing contact (name, email)	Customer order form	Contract performance	Duration of customer relationship + 7 years
License JWT issuance audit log (issuance ID, hash of issued JWT, issued-at, issued-by)	License-server internal	Legitimate interest — license enforcement, fraud prevention (Art. 6(1)(f))	7 years from issuance

The license-server runs in Gandiva-operated infrastructure. See **Section 7** for the security posture.

3.2 Support data

Purpose: to respond to customer support requests, debug deployments, and triage bugs.

Field	Source	Lawful basis (GDPR)	Retention
Support ticket content (text, attachments uploaded by customer)	Customer-initiated	Contract performance	90 days post-ticket-closure
Diagnostic bundles, log excerpts, screenshots (only when voluntarily attached by customer)	Customer-initiated	Contract performance	90 days post-ticket-closure

Support data is held in our ticketing system. Diagnostic bundles may include redacted operational logs from the customer's deployment; customers control what they choose to share.

3.3 Marketing-site data

Purpose: to operate the public marketing site at srasta.ai.

Field	Source	Lawful basis (GDPR)	Retention
Analytics consent choice	Visitor selection in the website banner	Legitimate interest (Art. 6(1)(f)) — remember privacy choice	Until browser storage is cleared
Anonymous web analytics (page views, page title, referrer, country, UTM campaign fields, CTA clicks, pitch-deck opens/downloads, booking clicks, anonymous browser ID)	First-party analytics after visitor consent	Consent (Art. 6(1)(a)) / legitimate interest where permitted	13 months
Contact-form, pilot-request, feature-request, product-update, and security-feature-request submissions (name, email, company, role, message, selected intent)	Customer-initiated	Consent (Art. 6(1)(a)) / Contract performance	Until consent withdrawn or 24 months after last interaction, whichever is earlier, unless legal/accounting/security retention requires longer
Newsletter subscriptions (email)	Customer-initiated	Consent	Until unsubscribed
Known investor pitch interactions (email-token link opened, deck downloaded, booking click)	Investor outreach links and explicit deck requests	Legitimate interest (Art. 6(1)(f)) / consent where required	24 months after last interaction unless unsubscribed or deleted earlier

The marketing site uses a first-party consent banner before creating an anonymous analytics identifier or persisting campaign attribution in browser storage. It does **not** use third-party advertising cookies, cross-site behavioral tracking, or fingerprint visitors. We do not store raw IP addresses for marketing-site analytics; the edge capture path stores a daily rotating requester hash derived from request metadata for deduplication and abuse defense.

3.3.1 Sales and CRM use

Anonymous website activity is used as aggregated or anonymous intent signal. Anonymous visitors are not converted into fake CRM People records. If a visitor later submits a form, requests a pilot, subscribes to product updates, or follows a pitch-deck link tied to their email, we may associate the known contact with their request, campaign attribution, and related deck or booking interactions.

3.3.2 Account-level enrichment

Gandiva may evaluate account-level visitor identification or enrichment tools to understand company-level interest and sales fit. Before enabling any enrichment provider on srasta.ai, we will update this policy, document the provider in the sub-processor list where applicable, and maintain a privacy posture that avoids selling personal information or creating individual contact records from anonymous traffic without a lawful basis.

3.4 Operational telemetry

Purpose: to understand adoption, reliability, and engagement across Srasta installs. Telemetry is **on by default** and disabled per Section 3.4.4 below.

3.4.1 What's collected

The Srasta platform sends a daily HTTPS POST to a Gandiva- operated endpoint. Each payload is tied to the install's license_id and contains **counts and shapes only**:

Field class	Examples	Why we collect it
Install events	install start / success / failure timestamps	Funnel drop-off detection
Platform fingerprint	platform version, OS, arch, Python version, GPU model family	Compatibility surface; release planning
Cluster shape	node count, GPU count	Validates "1 node to N nodes" positioning; community-vs-enterprise edition signal
Activity flag	one bit per UTC day if any inference / tool call / admin action occurred	DAU / WAU funnel signal
Feature touches	counts of admin pages viewed, audit-feed loads, model-RBAC edits	Which features are used (count of touches; never what was viewed)
Engagement intensity	inference call count, tool call count	Curious-vs-serious signal
Error class counts	counts per error category (NOT error message content)	Stability + crash-frequency

3.4.2 What's NEVER collected

The telemetry payload is filtered to exclude — even by accident — any customer-controlled data:

- ❌ Prompt content, model response content, conversation history
- ❌ Document content from the customer's knowledge base
- ❌ Audit-log content (only error class **counts**, never bodies)
- ❌ End-user identities, names, email addresses, IPs, hostnames
- ❌ Internal IP addresses, hostnames, network topology
- ❌ Customer organization-specific payloads of any kind

Telemetry is metadata about the platform, never about what flows through the platform.

3.4.3 Lawful basis (GDPR)

Edition	Lawful basis	Retention
Community (free trial)	Contract performance (Art. 6(1)(b)) — telemetry is part of the no-fee evaluation arrangement; explicit at trial sign-up	24 months aggregated, 90 days raw
Enterprise (paid customer, default ON)	Legitimate interest (Art. 6(1)(f)) — operational improvement, balanced against minimal-payload + opt-out availability	12 months aggregated, 30 days raw
Enterprise with telemetry disabled	N/A — no data collected	N/A

3.4.4 Opt-out

- **Community + Enterprise editions**: set `SRASTA_TELEMETRY=off` in the Srasta install's environment. No payload leaves the install.
- **Enterprise edition**: also disable via the License page in the admin UI. Toggle is persistent in `platform_config`.
- **At install time**: the install wizard (web UI + headless) shows a clear consent screen; opt-out is a single-click.

Disabling telemetry does **not** affect platform functionality. A telemetry-disabled install is fully usable; we just lose visibility into engagement.

3.5 What Gandiva does NOT process

To make the boundary explicit:

- We do not process the **content** of prompts, responses, or documents that flow through customer Srasta deployments.
- We do not process the **identity** of customer end-users (analysts, developers, etc. who use Srasta to do their jobs) — that is mediated by the customer's IdP (Zitadel / Okta / Azure AD).
- We do not process the customer's **audit logs**. Those live in the customer's deployment and ship only to customer-configured sinks. Operational telemetry (Section 3.4) collects **counts** of error classes, never the audit-log content itself.

---

4. Customer's Role and Our Role

Under GDPR terminology:

- For the personal data the **customer's end-users** generate inside the Srasta deployment, the **customer is the controller** and Gandiva is **not a processor**. Gandiva does not have access to that data.
- For the personal data **Gandiva itself processes** as listed in Section 3 (license, support, marketing), **Gandiva is the controller** for the marketing-site collection and **a joint controller / processor** for license + support data depending on contract terms.

The standard Srasta DPA (`docs/legal/msa-template.md` Annex / future `docs/legal/dpa-template.md`) codifies this allocation.

---

5. International Transfers

Gandiva-operated infrastructure (license-server, support ticketing, marketing site) is hosted in the United States.

For EU/UK/Swiss customers, transfers of license-server and support data to the US rely on:

- **Standard Contractual Clauses (SCCs)** — Module 1 (controller- to-controller) for marketing data; Module 2 (controller-to- processor) for support and license-server data, incorporated into the DPA.
- A transfer-impact assessment (TIA) addressing US government access risk.

We do not transfer customer end-user data internationally because we do not receive it. Where the customer's deployment runs is the customer's decision.

---

6. Sub-processors

Gandiva uses a small set of sub-processors strictly limited to operating the surfaces in Section 3. None of them have access to customer end-user data because that data never reaches Gandiva.

Current sub-processor list:

Sub-processor	Purpose	Region	Customer data exposure
AWS / GCP (cloud provider TBD)	License-server hosting, marketing-site hosting	US	Section 3.1 + 3.3 only
Linear (or similar)	Support-ticket workflow	US	Section 3.2 only
Stripe (when paid billing is enabled)	Billing	US	Billing contact + invoice data
Postmark / Resend (transactional email)	License delivery + support replies	US	Email addresses + content of emails we send

Sub-processors are listed in DPA Annex 2 and are kept current as the operating stack evolves. Material additions are notified to customers in advance per the DPA notification clause.

---

7. Security Posture for Gandiva-Held Data

Gandiva-operated infrastructure (license-server, support, marketing) is run with the following baseline:

- **Encryption at rest** — AES-256 via the cloud provider's managed keys.
- **Encryption in transit** — TLS 1.2+ for every external endpoint; mutual-TLS or signed JWTs for service-to-service.
- **Access control** — least-privilege IAM; production access requires a documented justification + 2-person review for long-lived sessions.
- **Audit** — every administrative access to license-server data is logged.
- **Key rotation** — quarterly for service credentials; immediate on suspected compromise.
- **Backups** — license-server database is snapshotted hourly, retained 30 days, encrypted.

- **Vulnerability management** — image CVE scanning in CI; critical-severity patches within 48 hours of confirmed CVE.
- **Incident response** — 72-hour notification per GDPR Art. 33 to affected customers + supervisory authority where required.

For the controls matrix mapping these to SOC 2 Common Criteria, see `docs/security/controls-matrix.md`.

---

8. Your Rights (GDPR / UK GDPR / CCPA)

If you are a EU/UK/Swiss data subject or a California resident, you have the following rights with respect to personal data Gandiva controls (Section 3):

- **Access** — request a copy of the personal data we hold about you.
- **Rectification** — correct inaccurate data.
- **Erasure** — request deletion (subject to retention obligations for tax/accounting and dispute-resolution).
- **Restriction** — limit how we process your data while a dispute is resolved.
- **Portability** — receive your data in a structured, machine- readable format.
- **Objection** — object to processing based on legitimate interest.
- **Opt-out of "sale" / "sharing"** (CCPA) — not applicable, since we do not sell or share personal data.
- **Lodge a complaint** with a supervisory authority (your home data-protection authority for GDPR; the California Attorney General for CCPA).

Submit requests to **privacy@srasta.ai** with sufficient detail to verify your identity. We will respond within 30 days (extensible to 60 days for complex requests, with notice).

For data inside a customer's Srasta deployment, requests must be addressed to the customer (the controller), not to us.

---

9. Children's Data

Srasta is a B2B platform. Gandiva does not knowingly process personal data from individuals under 16. If you become aware that a child has submitted personal data to us, contact privacy@srasta.ai.

---

10. Cookies and Similar Technologies

The marketing site at srasta.ai uses first-party analytics only after the visitor accepts analytics in the privacy banner. The banner stores a consent choice in browser storage. If accepted, the site stores a randomly generated anonymous browser ID and short-lived campaign attribution in browser storage so we can measure page visits, CTA clicks, pitch-deck interactions, and pilot interest. Declining analytics does not block access to the site.

Srasta does not use third-party advertising cookies, cross-site behavioral tracking, or fingerprinting on the marketing site.

The Srasta platform itself uses session cookies inside the customer's deployment for OIDC login; those cookies are governed by the customer's privacy policy, not this one.

---

11. Changes to This Policy

We will post material changes to this policy at https://srasta.ai/privacy and notify enterprise customers via the DPA-stipulated change channel. The "Last Updated" date at the top of this document reflects the most recent revision.

---

12. Contacting Us

- **Privacy questions / data subject requests:** privacy@srasta.ai
- **Security disclosures:** security@srasta.ai (see `docs/security/responsible-disclosure.md` for details — planned)
- **General inquiries:** info@srasta.ai
- **Postal:** Gandiva Tech, Inc. — [Address TBD post-seed-close]

Related documents

- `docs/security/architecture.md` — the customer-perimeter architecture this policy is anchored on.
- `docs/security/controls-matrix.md` — SOC 2 Common Criteria mapping covering the controls referenced in Section 7.
- `docs/security/caiq-lite.md` — CAIQ Lite questionnaire responses that mirror this policy's data-handling assertions.
- `docs/legal/dpa-template.md` — Data Processing Agreement template (planned via #172) — codifies the controller / processor allocation in Section 4.
- `docs/legal/msa-template.md` — Master Services Agreement template (planned via #172).
- `docs/strategy/seed-momentum-plan.md` — closure-boundary item #6 (privacy policy + CAIQ Lite).