DataSitr — Security & Privacy Summary

Generated: 2026-05-17T05:18:20Z

Source document: docs/customer-security-one-pager.md

Git commit: 460842ce

Generator: [operator-tooling]

Benchmark artifact: docs/generated/pii_benchmark_latest.json (generated 2026-04-29T14:40:52Z, Arabic NER loaded, 1K p95 47.92 ms)

For: Customer security review, procurement, vendor assessment

---

What DataSitr does

DataSitr is a privacy gateway that lets Saudi organizations use external AI services without sending raw personal data to cross-border AI providers. It detects PII, tokenizes it before external transmission, and restores original values from an encrypted in-Kingdom vault after the AI responds.

Core privacy design

Raw personal data is not sent to cross-border AI providers by design. External AI providers only receive detector-sanitized text after tokenization and post-tokenization rescan. Amber/red handling is intended to remain on operator-configured in-Kingdom paths, which DataSitr does not independently verify. Residual contextual re-identification risk and legal reliance still require customer and legal review.

How data is protected

Layer	Mechanism
Data classification	Automated PII detection (names, National IDs, Iqama, phone, IBAN, medical, biometric) using NLP + Saudi-specific pattern recognizers
Three-lane routing	Green (tokenized, external AI) / Amber (pseudonymized, operator-configured in-Kingdom path) / Red (raw, strict in-Kingdom path) / Block (rejected)
Tokenization verification	Post-tokenization rescan confirms no original PII remains before external transmission

Sensitive data handling	PDPL Article 1(11) categories (health, biometric, genetic, criminal) are routed to the strictest in-Kingdom path available — never sent to cross-border providers
Encryption at rest	AES-256-GCM with per-tenant derived keys; random nonce per operation
Encryption in transit	TLS 1.2/1.3 with HSTS enforced
Tenant isolation	Cryptographic, access-control, logging, policy, and rate-limit isolation per tenant

Authentication & access control

Feature	Detail
API authentication	Bearer token with SHA-256 hashing and constant-time comparison
Role model	`super_admin` / `tenant_admin` / `tenant`, plus a separate read-only `regulator` role when provisioned — enforced at API and dashboard layers
SSO integration	OIDC (Authorization Code + PKCE) supporting Keycloak, Microsoft Entra ID, Okta, Auth0
Session security	HttpOnly, Secure, SameSite=Lax cookies; HMAC-SHA256 signed session IDs

Compliance posture

Requirement	Implementation
PDPL cross-border transfer controls	Three-lane routing with fail-closed behavior
Processing records	Machine-readable append-only records per SDAIA 5-year retention guidelines, with hash-chain continuity and keyed HMAC authentication for records created after the integrity feature deployment (legacy records lack HMAC)
Transfer register	Every routing decision is recorded, with cross-border transfer rows explicitly distinguished by transfer type, decision rationale, provider, and policy version
Subject rights	Data export, deletion, and audit trail implemented and tested
Article 1(11) sensitive data	Classified and routed to in-Kingdom processing only

For the operator/legal handoff package behind cross-border green-lane review,

use Transfer Governance Package.

Operational security

Area	Status
Test coverage	A dated verification snapshot is maintained separately and refreshed before buyer-facing counts are reused; covered surfaces include auth boundaries, tenant isolation, webhook durability, monitor health, deploy / backup / restore scripts, and dashboard auth flows
Deployment	Live pilot runs on a Saudi-hosted shared-state stack; guarded VPS and guarded ACK/Helm deployment paths both exist in-repo
Backup/restore	Latest dated pilot evidence covers encrypted-first backups, verified OSS upload/download, and restore-drill execution; the newest encrypted archive is the active off-host candidate, monitor coverage warns on plaintext drift, and SSH fallback remains available. Continuity evidence is operator-directed and backup-based, not replication-backed or automatic
Monitoring	Health endpoint probes, alert webhook delivery, provider circuit breakers, and the pilot monitoring/metrics stack. Alert freshness is proved by dated operator evidence, not by this summary alone
Continuity evidence	Dated proofs span March 24–29 (single-host `worker=2`, first two-instance and two-host shared-state, planned rolling deploy, isolated restore recovery, public OIDC failover, public restored-state cutover from a dated encrypted backup, public Host A ingress independence from restored local state), the 2026-04-21 guarded rollout baseline, the 2026-05-04 multi-AZ ACK customer-route cutover with 4-hour soak (signed bundle at `evidence/ha/alibaba-live-2026-05-04T01:17:03Z/`), and the 2026-05-16 GCP Dammam warm-standby drill rehearsal (Ed25519-signed bundle at `evidence/multi-region-drill/multi-region-warm-standby-20260516T220433Z/`, scope: DNS / GKE Ingress / TLS routing / Cloud Armor WAF only). These are separate proof boundaries, not cross-cloud replication-backed continuity, automatic database/auth failover, full-vault verification, unplanned full Host A machine-loss tolerance, unplanned full-region tolerance, or blanket HA
Service hardening	systemd `NoNewPrivileges`, `ProtectSystem=strict`, non-root execution
Incident response	Documented procedures for PII leak, provider outage, vault corruption, DDoS
Key rotation	Master key rotation with dry-run; runtime secret rotation scripted

<!-- claims-sync:customer-security:start -->

Claim Boundary Snapshot

Generated from `docs/claims_manifest.json`. This section states the current buyer-safe claim boundary for security review material. Update `docs/claims_manifest.json`, then rerun `python3 [operator-tooling].`

Safe To Say Now

- DataSitr is in pilot posture and is not being represented as blanket HA, regulator-approved, or enterprise-certified.
- Quasi-identifier detection and structured quasi-risk assessment are implemented and enforced in anonymization checks.
- Consent-basis processing requires an explicit subject_identifier, and the pipeline blocks matching withdrawn-consent requests before tokenization or provider work.

Say Only With Explicit Caveat

- Immutable-evidence export tooling exists, but PostgreSQL/shared-state evidence does not justify a universal immutable-evidence claim. Caveat: Use careful wording such as export tooling, tamper-evident chains, or bounded evidence proofs. Avoid saying the whole system is immutable.
- Tenant BYOK is launch-ready for the narrow v1 Alibaba KMS provider-credential path. Caveat: Do not say a named tenant is live under BYOK until the evidence gate proves runtime config, custody canary, stored wrapped credential, provider-auth validation, runtime unwrap audit, and fail-closed behavior. Do not claim HSM-backed custody.

See Claims Registry for the full generated registry.

<!-- claims-sync:customer-security:end -->

Current limitations (disclosed transparently)

Area	Status	Mitigation path
Independent security audit	Not yet performed	Planned before broader commercial rollout
HSM / external KMS	Runtime still depends on operator-managed secret material at startup/runtime; startup KMS bootstrap exists, but external KMS/HSM is not the steady-state live custody claim	HSM/KMS integration and custody hardening remain separate rollout steps
Cryptographic log signing	SHA-256 hash chain plus keyed HMAC authentication for new records; PostgreSQL ordering depends on sequencing metadata; signed evidence export (Ed25519) exists for configured sequenced rows. Compliance-record HMAC can use an independently configured key (`SV_COMPLIANCE_HMAC_KEY`), with backward-compatible fallback to master-key derivation if unset; it is still not an independent external signature. Legacy records created before the HMAC feature lack keyed authentication.	Keep the signed-evidence OSS path current; WORM lock and universal legacy-row coverage remain separate rollout steps

Business continuity boundary	Current restore and public continuity proofs are operator-directed and backup-based	Replication-backed continuity, automatic failover, and broader HA remain separate engineering and proof steps
Arabic NER accuracy	Benchmark and evaluation artifacts exist, but tenant-specific workload validation is still required	Low-confidence cases route fail-safe to in-Kingdom AI; keep buyer claims tied to current benchmark artifacts rather than blanket accuracy claims
Legal precedent	PDPL anonymization exemption untested in Saudi courts	Safety-first routing; legal counsel recommended

Certifications

DataSitr does not currently hold SOC 2, ISO 27001, or equivalent certifications. The architecture is designed to support future certification efforts. An independent penetration test and security review are planned.

For the current reviewer-safe immutable-evidence packet, use

Immutable Evidence Handoff Package.

Version: 0.1.1 | **Last updated:** 2026-05-17

For the canonical list of safe and unsafe claims, see Claims Boundary.

See also: Security & Compliance Overview | Transfer Governance Package | Immutable Evidence Handoff Package | Threat Model | Tenant Isolation

This document describes technical design intent and current operational posture. It does not constitute a warranty, service-level agreement, legal guarantee, or certification of regulatory compliance. DataSitr is designed to support PDPL alignment; it does not itself grant compliance. For the canonical list of safe and unsafe claims, contact gov@datasitr.com.