DataSitr — Customer Evaluation Guide

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

Source document: docs/customer-evaluation-guide.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: Technical champions, security reviewers, pilot customers

This guide walks you through evaluating DataSitr from first contact to pilot sign-off.

What you are evaluating

DataSitr is a privacy gateway that supports Saudi PDPL alignment when your applications use external AI services. You are evaluating whether it:

1. Detects the PII your business handles (Saudi IDs, phone numbers, IBANs, names, etc.)
2. Routes requests through the correct privacy lane for your data sensitivity
3. Produces the compliance records your legal and compliance teams need
4. Integrates with your existing applications without workflow changes
5. Meets your organization's security and procurement requirements

---

Evaluation timeline

Phase	Duration	What happens
Technical intro	1 day	Demo call, review architecture, answer security questionnaire
Pilot setup	1-2 days	Provision the Saudi-hosted pilot stack or customer-hosted deployment, issue API keys, verify health
Integration test	1-2 weeks	Connect your application, process representative data, review dashboard
Security review	Parallel	Security team reviews documentation pack (see below)

Pilot sign-off	1 day	Review results, agree on production terms or extend pilot

Total: 2-4 weeks for a typical evaluation.

Step 1: Security and compliance review

Before integration, your security team will want to review DataSitr's posture. We provide a pre-built evidence pack:

Document	What it covers
Customer Security One-Pager	2-page executive summary of the privacy model, encryption, auth, and disclosed limitations
Security Questionnaire Answer Pack	Pre-drafted answers to common vendor security assessment questions
Compliance Integrity Handoff Package	Reviewer-safe packet for the current processing-record verifier story, synthetic fixtures, and exact integrity claim boundary
Transfer Governance Package	Repo-safe legal/privacy handoff pack plus fillable templates for provider inventory, TIA, approval, and disclosure alignment
Immutable Evidence Handoff Package	Reviewer-safe packet for the current signed-evidence verifier story and proof boundary
Compliance Reviewer Pack - 2026-05-17	Current reviewer-facing packet summarizing the 144-control matrix, signed bundle, detector benchmark artifacts, and explicit claims boundary
Backup Hardening Summary - 2026-03-30	Dated note on the encrypted-first backup posture, verified off-host state, and exact continuity non-claims
Production Readiness Checklist	Current readiness checklist, each item traceable to tests or scripts
Disaster Recovery Summary	Backup/restore posture, failure scenarios, recovery procedures
Threat Model	Threat catalog, mitigations, residual risk assessment
Tenant Isolation	How tenants are cryptographically and operationally isolated
Architecture Overview	Request flow diagram, module structure, data residency map

We disclose known gaps honestly. The continuity boundary today is encrypted, operator-directed, backup-based restored-state continuity, not replication, automatic failover, or blanket HA.

If your security team uses an independent reviewer, start them with the current Compliance Reviewer Pack - 2026-05-17.

The dated

Third-Party Diligence Scope - 2026-03-30

is retained as a historical appendix only; do not use it as the current

reviewer entry point.

If your privacy/legal team needs the operator-side transfer packet rather than

just the product docs, send them

Transfer Governance Package.

Step 2: Pilot deployment

DataSitr is deployed on a Saudi-hosted stack. The current public pilot uses a

Saudi-hosted shared-state deployment, and the repo contains guarded VPS and

guarded ACK/Helm rollout paths. Raw personal data does not leave Saudi Arabia

by design. Green-lane external transfer is limited to detector-sanitized text

after tokenization and post-tokenization rescan; residual contextual

re-identification risk still requires customer and legal review.

Deployment options:

Option	Best for
DataSitr-hosted pilot	Fastest start — we provision and manage the Saudi-hosted stack
Customer-hosted VPS/systemd	Your team deploys on your own Saudi VPS or datacenter host using our deployment guide
Operator-assisted shared-state / Helm path	Evaluations that need closer shared-state or production-parity posture; handled as a guarded rollout, not a self-serve quickstart

Either way, you get:

- Full three-lane routing (green, amber, red)
- Customer dashboard at `https://<your-instance>/dashboard/`
- API access at `https://<your-instance>/v1/`
- TLS encryption in transit (HSTS enabled)

See Deployment Options for details on both paths.

Step 3: Integration

Integration is a single API endpoint. No SDK is required.

[See source document for diagram/code]

The API Quickstart gets you from zero to first API call in under 5 minutes.

What to test during integration:

Test	How
PII detection accuracy	Send representative text through `/v1/detect` and review entity types and confidence scores
Lane routing	Process requests and verify lane assignments match your expectations in the dashboard
Compliance records	Check the Compliance tab for processing records and transfer register entries
Subject rights	Test data export and deletion for sample subject identifiers
Webhook delivery	Configure a webhook URL and verify signed delivery
Force in-Kingdom	Set `"force_in_kingdom": true` and confirm all processing stays local
Error handling	Test with invalid keys, rate limit boundaries, and provider failures

---

Step 4: Dashboard review

The customer dashboard provides visibility into everything DataSitr does with your data:

Tab	What you see
Overview	Health status, usage stats, lane distribution, provider status
Compliance	Processing records, transfer register, RoPA export, DPIA
History	Per-request detail: text (redacted), lane, provider, timing, cost
Billing	Statements, rate-card visibility, and tenant billing context
Subject Rights	Search, export, delete, and rectify subject data
Settings	API keys, tenant controls, and role-scoped administrative settings
Academy	Guided product and compliance training by role
Help	Searchable articles and interactive API reference

Tab visibility is role-based. See Dashboard Help for a full walkthrough.

Step 5: Pilot sign-off

At the end of the evaluation, review:

Criterion	Evidence
PII detection covers your data patterns	`/v1/detect` results, dashboard Compliance tab
Lane routing matches your compliance requirements	Processing records, transfer register
Integration effort is acceptable	Time to first API call, code changes needed
Security posture meets your requirements	Evidence pack review, security questionnaire answers
Billing model works for your use case	Dashboard billing statements, rate card discussion

Possible outcomes:

- **Proceed to production** — Agree on rate card, provisioning timeline, and support terms
- **Extend pilot** — More testing time, broader data coverage, additional integrations
- **No-go** — Clean shutdown, all data deleted per PDPL subject rights

---

What you will need

Item	When	Notes
Representative text samples	Integration phase	Real or realistic data your application sends to AI providers
AI provider API keys	Setup	At least one external (OpenAI/Anthropic/Google) and one in-Kingdom (STC SambaNova) for full lane coverage
Security questionnaire (if applicable)	Review phase	We provide pre-drafted answers — see Security Answer Pack
Webhook endpoint (optional)	Integration phase	HTTPS URL for async result delivery
Corporate IdP details (optional)	If SSO is needed	OIDC-compatible IdP (Entra ID, Okta, Auth0) — see OIDC SSO Guide

---

Frequently asked questions (evaluation-specific)

Can I test without real customer data?

Yes. DataSitr detects synthetic Saudi data patterns the same way it detects real ones. Use realistic test data during integration, then validate with real data (with consent) when ready.

What if my data includes Arabic text?

DataSitr detects Arabic text with Saudi patterns, an Arabic name dictionary, and a configurable local Arabic semantic backend enabled by default. Pilot feedback on Arabic accuracy still helps tune recall and precision for customer-specific text.

What happens to my data after the pilot?

Vault entries auto-expire (default 24 hours). Vault-linked subject data can be

exported, rectified, deleted, or marked as consent-withdrawn through the

subject-rights flows. Retained compliance records follow the 5-year SDAIA

retention posture and keep the corresponding audit markers instead of being

silently erased.

Is there a free tier?

The pilot phase is free or negotiated. Production pricing is per-request. See Pricing & Packaging.

Related docs

- API Quickstart — First API call in 5 minutes
- Deployment Options — Hosting and deployment paths
- Pricing & Packaging — Billing model and tier structure
- Enterprise Readiness FAQ — SSO, certifications, SLAs, multi-tenancy
- Transfer Governance Package — Cross-border legal/privacy handoff package
- Immutable Evidence Handoff Package — Signed-evidence reviewer handoff package
- FAQ — General product questions

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

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.