

Design Partner Brief

Sakshi · AI agent governance | For regulated Indian BFSI | July 2026

AI agents are moving into real decisions in Indian banking, credit, cards, and insurance: underwriting, fraud screening, claims triage, suitability. The regulators are moving with them. This brief is for a BFSI team evaluating Sakshi as a design partner. It is meant to be read by the people who sign off: security, risk, compliance, and the sponsor who owns the use case.

Category	Agent governance	Built for	Regulated Indian BFSI
Tenancy	Single-tenant, your VPC	Raw PII stored	None
Regulatory maps	RBI, DPDP, SEBI, IRDAI	Evidence	Hash-chained, signed
Ways in	Sandbox to production	Components	Open, Apache-2.0

01 Why now

The rules are arriving on a clock, and they name specific controls.

- **RBI.** The draft Model Risk Management directions name an inventory of models and agents, a kill switch, AI disclosure with human handoff, and red-teaming as expectations. Final directions are expected in H2 2026.
- **SEBI.** Regulation 16C is already binding, and it puts responsibility for AI systems, including third-party ones, on the regulated entity.
- **IRDAI and DPDP.** An IRDAI AI framework is anticipated around the second half of 2026, and DPDP significant-data-fiduciary duties (algorithmic due diligence, DPIA) follow. The obligations are converging on the same evidence.

The window is to have the governance layer in place before an auditor or a board asks for it. No global platform ships these India-specific mappings; the hyperscalers govern their own stacks, not yours.

02 What Sakshi does

- **Register.** A live inventory of every agent and model, with an owner, an autonomy tier, provenance, and a verifiable identity. Nothing runs unregistered.
- **Witness.** Every decision is hashed onto a tamper-evident chain with the model identity, tokenized inputs, and outcome, recomputable by anyone.
- **Bound.** Autonomy envelopes route each action to auto, human review, or block by stakes and confidence, with oversight telemetry and a drilled kill switch.
- **Vidhi and Vishwas.** Regulator instruments encoded as data and scored against live evidence, plus declared-first fairness screens gated on statistical significance.

03 What it produces for the regulator

Sakshi is built to produce regulator-grade artifacts from live evidence, not slideware.

INSTRUMENT	WHAT SAKSHI PRODUCES
RBI MRM (draft)	Model and agent inventory, decision evidence, drilled kill switch with signed attestations, incident register aligned to the FREE-AI annexure, red-team evidence.
DPDP	PII tokenization, data-minimization and localization screens, and an independent-auditor DPIA and audit bundle.
SEBI Reg 16C	Third-party model attribution on the chain, AI disclosure, and suitability decision evidence.
IRDAI (anticipated)	A skeleton clause map that re-cuts on the final report, riding the same evidence.

Each instrument gets a readiness score computed from live checks, and a signed evidence pack that is Ed25519-signed over its exact bytes. The mappings are data, so a new instrument is a new file, not a new build.

04 Security and deployment posture

- **Single-tenant, in your environment.** Each customer runs their own instance in their own VPC or on-prem. No shared control plane, no call-home, works air-gapped.
- **Raw PII never stored.** Identifiers are tokenized at ingest, before anything is written or hashed. The signing seed and tokenization config stay in your instance; Rotavision cannot detokenize your records or forge a signature.
- **Independently verifiable.** Recompute the chain to re-derive every hash; a single changed field is located exactly. Your teams and your regulator do not take our word for it.
- **Identity and separation of duties.** SSO via OIDC or SAML against your IdP, role-based access with a read-only auditor role, and entitlement changes attested on the chain.

On formal certifications we are deliberately plain: SOC 2 Type II and ISO 27001 are on our roadmap, not yet in hand, and we will not imply them before they land. The components you run are open and Apache-2.0 licensed, so they are inspectable rather than taken on trust. A full Security Overview is available at rotavision.com/security.

05 What a design-partnership includes

- **We deploy alongside you.** Single-tenant in your VPC or on-prem, wired to the agents you actually run, with a working install in an afternoon.
- **We shape the roadmap around you.** Priority is set by what you hit first, deepest in BFSI where the regulatory work is furthest along.
- **Named support and regulatory advisory.** Direct access to the team, and help mapping your use case to the instruments that apply to it.
- **What we ask.** A real use case, an executive sponsor, and honest feedback. The sandbox and the open SDKs are free to start.

06 How to evaluate in about thirty minutes

- **Get a writable instance.** Request a private synthetic-data sandbox at rotavision.com/sandbox, or self-host with one command in your own VPC.
- **Govern your first agent.** A few lines with the SDK, or one line of middleware for OpenAI, Anthropic, Gemini, Bedrock, LangGraph, Google ADK, or MCP.
- **Verify and export.** Recompute the chain yourself, then export a signed, regulator-mapped evidence pack. The full walkthrough is at docs.rotavision.com/evaluate, and the deployment and posture detail at docs.rotavision.com/self-hosting.

Next step. Tell us the use case and we will scope a design partnership with you, or open a sandbox now and try it first. A completed security questionnaire and a data processing addendum are available under NDA. Reach us at hello@rotavision.com, or rotavision.com/contact.

Rotavision Consulting Private Limited · AI Trust Infrastructure for India

WeWork, HD-37, Block D3, Manyata Tech Park, Bengaluru 560045, Karnataka, India

rotavision.com · hello@rotavision.com · security@rotavision.com