Sahamati · Google Cloud

SAHAMATI LABS · GOOGLE

Aikya

The Fraud Intelligence Layer for India's Open Finance Ecosystem

A presentation by Kiran Gopinath, Sahamati Labs · Rene Kolga, Google Cloud

WHO BUILDS THIS

Sahamati Labs

The innovation arm of Sahamati Foundation, A Reserve Bank of India (RBI) recognized Self Regulatory Organization for India's Open Finance ecosystem. We build the next versions of Account Aggregator (AA) — India's Open Finance framework — extending, hardening, and improving what's deployed today.

SOME OF OUR PROJECTS · EACH WILL USE CONFIDENTIAL COMPUTING

Aikya

Fraud intelligence for the AA ecosystem. Cross-lender signals via confidential computing.

Cross-border data sharing

Interoperability between Open Finance jurisdictions, portable financial signals across borders.

Agentic AI guardrails

Governance framework for AI agents operating on consent in the AA ecosystem.

Open Data

Bringing business and commercial data sources onto AA, the first move toward Open Data.

Today's talk: **a deep dive into Aikya.**

THE LANDSCAPE

India's Open Finance has reached ecosystem scale

AA is the connective tissue of digital credit in India. Borrowers can trigger multiple consented data pulls instantly but every lender still sees only their slice.

989

FIUs

data-consuming institutions

179

FIPs

data-providing institutions

285M

accounts linked

consumer adoption to date

431M

consents fulfilled

data-sharing transactions

Speed without intelligence increases ecosystem vulnerability. *Lenders need real-time visibility.*

Stacked to a few million. Then gone.

Recounted by a large lending institution. Anonymised, but representative of how stacking happens.

HOW THE CASE UNFOLDED

Stacked across lenders

Drew loans from several institutions in parallel. Each approved on the slice it could see; none saw the others.

A few million, in aggregate

Exposure compounded across lenders into the low millions of dollars before any bureau update caught up.

Then vanished

By the time the lenders moved to trace the borrower, they had already left the country.

WHERE AIKYA FIRES

The full picture only ever existed across the ecosystem.

Aikya computes the cross-lender stacking signal in the same hop as the AA pull, surfacing the aggregate exposure before disbursement, not after the borrower is gone.

Catch the stack at the application, not at the airport.

THE PROBLEM

AA fraud is rising — and invisible to any one lender

HOW BORROWERS EXPLOIT BLIND SPOTS

Parallel applications

Apply across multiple lenders simultaneously, before any one updates the bureau.

Synthetic & mixed identities

Fabricate or blend identifiers to evade per-lender deduplication.

Build-then-bust-out

Establish good credit, then draw the maximum across lenders in one shot.

USD 4.06B

lost to fraud in 2024–25

across India's regulated financial sector · ~₹36,000+ crore

92%

of fraud losses were loan-related.

Personal loan NPAs rose 51% and credit card NPAs 136% over the same period.

THE OPPORTUNITY

What the AA ecosystem needs

Five capabilities the ecosystem is missing today — and that no single lender can build alone.

1. Real-time fraud detection
2. Cross-lender behavioural insight
3. Early warning of risky borrowers
4. Consent-safe intelligence
5. Zero data sharing

Inside the enclave: end-to-end

01 Attestation

Before dispatching data, the Financial Information User (FIU) validates the enclave's signed attestation and measurement against the approved Aikya enclave build, ensuring that data is encrypted only to a trusted runtime.

02 Encrypted ingress

Query parameters and AA-derived counts are encrypted with a key the enclave generates and seals to itself. The key is never exported. No operator, including Sahamati, sees plaintext.

03 Sealed compute

Pattern detection runs inside hardware-isolated memory. The host OS, hypervisor and cloud provider cannot read inputs or intermediate state. Memory is encrypted by hardware.

04 Signal egress, state destroyed

Only the risk signal leaves the enclave. Working memory is wiped when the enclave shuts down. There is no archive of inputs to subpoena.

Verifiable from the outside, opaque from the inside. *Every FIU can prove what the network does without seeing what any other FIU sent.*

Aligned with AA's data-blind principle

AA does not permit a central aggregator of consent activity. Confidential computing is what lets Aikya compute ecosystem-wide signals without ever creating one.

1 No central data lake

Counts are computed on the fly inside the enclave and discarded after the signal is released. Nothing accumulates.

2 Sahamati cannot see

Even when compelled by a court order, there is no plaintext to disclose. The operator is excluded from the trust boundary.

3 Verifiable code, not promises

Every release of the enclave binary is published and attestable. Trust comes from cryptography, not from a privacy policy.

4 Aligned with the consent model

Each input is bound to an active consent artefact. No consent, no input. Revocation invalidates downstream computation.

INTRODUCING AIKYA

Aikya is the privacy-preserving, AI-enabled fraud intelligence network for the AA ecosystem.

Behavioural fraud signals

Cross-lender patterns derived from consent metadata.

Real-time risk visibility

Velocity and stress signals in the same hop as the AA pull.

Privacy by design

No raw data leaves any FIU. Counts and signals only.

Confidential computing

Hardware-isolated enclaves with verifiable code.

AI-driven detection

Anomaly and pattern detection over network-level features.

HOW IT WORKS

A query goes in. A signal comes out. No data moves.

01

Lender query

FIU sends an encrypted query to the Aikya network.

02

AA counts

AAs return anonymised, aggregated counts — never identifiers.

03

Enclave compute

Confidential computing encloses all processing inside attested code.

04

Signal out

Aikya returns a risk signal and pattern indicators to the lender.

No PII. · No raw data movement. · No workflow changes.

THE ARCHITECTURE

Federated query, encrypted compute

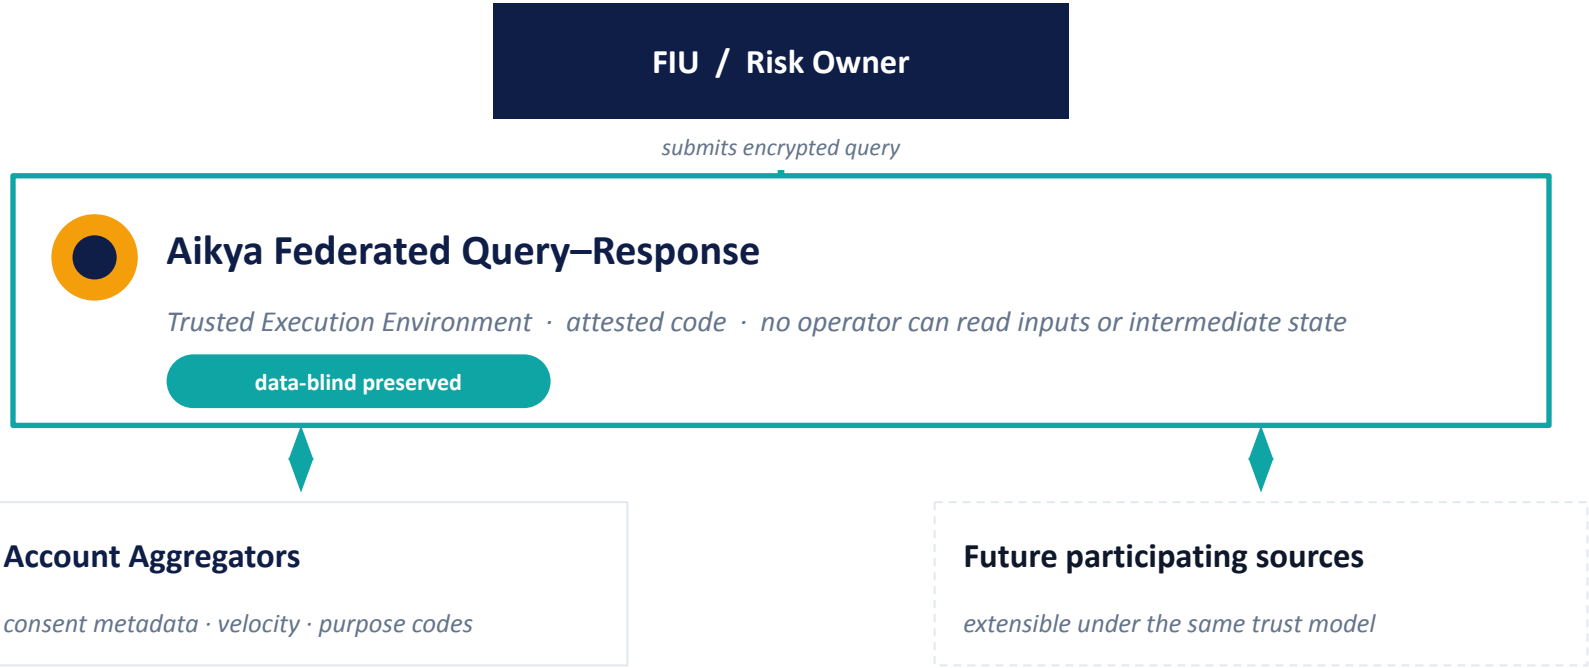

Output: *a risk signal — never raw consent data, never a row from any FIU's pull.*

PLATFORM

Built on Google Cloud Confidential Space

Confidential Space is Google Cloud's confidential computing platform for multi-party data collaboration. The pilot sandbox runs on Confidential Space today.

Hardware-backed enclaves

Workloads run inside Confidential VMs on AMD SEV or Intel TDX, memory is encrypted by the CPU.

Verifiable attestation

Every workload publishes signed OIDC attestation tokens that FIUs can verify before sending any data.

Multi-party by design

Purpose-built for collaboration where multiple parties contribute data to a shared computation without trusting one another.

Sealed key release

Workload Identity Pool gates access to Cloud KMS / HSM keys based on attestation claims, keys only released to verified enclave code.

CAPABILITIES

Aikya's core capabilities

01 Detect rapid loan applications

Real-time visibility of multiple AA pulls against the same identifier across lenders.

02 Understand borrower intent

Differentiate fresh applications from routine monitoring and account checks.

03 Lower false positives

Multi-signal analysis with behavioural context, AI pattern detection and custom lender thresholds.

04 Instant holistic risk snapshot

See how many lenders are checking, monitoring, or registering stress signals on a borrower right now.

WHERE WE ARE

Ready for pilots

Aikya is past the architecture stage. The core primitives are running; the next milestone is supervised production pilots with FIUs and AAs.

✓ **Test environment is live and accessible**	FIUs can already exercise the velocity-check endpoint with real-shaped traffic.
✓ **Velocity-check API functional**	End-to-end query → enclave → signal flow operating with attestation.
✓ **Lender POCs ready to onboard**	First pilot cohort identified; integration scope and SLAs scoped.
✓ **AI behavioural models in development**	Pattern, anomaly, and scoring models trained on representative consent metadata.
✓ **Strong interest across the AA ecosystem**	Engagement from AAs, FIUs and large lenders on shaping the governance layer.
○ **Awaiting regulator approval**	Production rollout is the final gate; supervised pilots will begin upon clearance.

Sahamati · Google Cloud

AIKYA · SAHAMATI LABS · GOOGLE

Thank you

The privacy-preserving, AI-driven fraud intelligence layer for India's Account Aggregator ecosystem.

Protecting lenders. Preserving trust. ***Scaling credit safely.***

Sahamati Labs · kiran.gopinath@sahamati.org.in