

Your integration layer is the one place your compliance boundary leaks.

Samba doesn't create that leak. The integration tier deploys **inside** the boundary you've already authorized — not in a vendor's tenant outside it.

WHY GOVERNMENT IS THE FIT

Defense Industrial Base

A CMMC Level 2 assessment forces an honest inventory of every system touching CUI. A multi-tenant integration platform in that inventory can be a disqualification, not just a finding. Samba lands inside the enclave you already have — same boundary, same IAM, same audit evidence — instead of adding a vendor to your SSP.

Federal Civilian

The customer-owned boundary model removes the FedRAMP vendor-authorization problem entirely. You aren't asking us to go get authorized. You're extending the authorization you already hold, to a system you already operate.

State, Local & Higher Ed

StateRAMP, FERPA, and CJIS obligations on a fraction of federal budget. No new vendor authorization to sponsor, no new data-sharing agreement to negotiate, and no additional third party in scope at renewal.

The common thread: **your authorization boundary already exists.** Samba deploys inside it, instead of asking you to extend trust to a new one.

BUILT INSIDE YOUR BOUNDARY

- Deploys into **your** Kubernetes cluster, under **your** IAM, on **your** network.
- Credentials, payloads, and audit logs never cross into a vendor tenant.
- FIPS-mode cryptography, and controls mapped to NIST 800-171 / 800-53 alongside SOC 2 — the paperwork your assessor already wants, in the format they already read.
- Source-code escrow available for enterprise contracts, so continuity risk isn't a reason to say no.


EVIDENCE, NOT A PROMISE

- **54 controls**, each mapped to a SOC 2 criterion and NIST 800-53 lineage, evaluated live against the cluster you're actually running.
- Daily backups, proven every week by performing an actual restore — not asserted annually by policy. Plus a database-enforced, tamper-proof audit log.
- Every configuration change, flow execution, and credential use logged with correlation IDs, user context, and old/new values.
- Your assessor reads it directly, in your environment, on their own schedule — nothing to take on faith.

THE SAME INTEGRATION, TWO ARCHITECTURES

SAMBA — SINGLE-TENANT, CUSTOMER-DEPLOYED


YOUR AUTHORIZATION BOUNDARY



Nothing crosses the line. The integration tier is one more workload inside a boundary your assessor has already signed off on.

MULTI-TENANT IPAAS

YOUR AUTHORIZATION BOUNDARY



↓ **credentials + payloads leave** ↓

VENDOR TENANT



A new boundary to authorize. Plus a vendor in your SSP, a risk assessment that renews annually, and a release schedule you don't control.

None of that matters if the platform doesn't do real work. Here's what it does.

CONNECT ANYTHING WITH AN API

- REST and HTTP endpoints — no pre-built connector library required, and none to wait on.
- XML and JSON sources alike, converted automatically for mapping.
- Pagination handled for you — next-link and offset cursors followed until the full set is retrieved.

MOVE DATA THE WAY THE WORK HAPPENS

- **Synchronize** — compare source and target, insert what's new, update what changed.
- **Push** — insert-only, when the target doesn't need reading.
- **Webhook ingestion** — inbound events processed in real time, routed to insert or update on a key match.
- **Composite matching** — match records on several key fields at once, not just one.
- **Incremental sync** — a timestamp parameter so the source returns only what changed.

THE AI HAS TWO JOBS — AND THE SAME HARD LIMIT ON BOTH

It helps you build integrations

- Describe the goal in plain language — the AI drafts the full integration plan.
- Identifies source systems, target systems, and enrichment lookups from context.
- Suggests method, URL pattern, and list path; parses pasted OpenAPI specs or sample payloads to extract fields.
- Writes field-level Python transforms, and refines conversationally as you revise.

HANDLE WHAT USUALLY MEANS CUSTOM CODE

- **Multi-step flows** — chain lookups and writes across several systems in a single run.
- **Source enrichment** — fetch related data from a second system before the target write.
- **Write-back** — stamp the target's returned ID back onto the source record.

OPERATE IT LIKE REGULATED SOFTWARE

- Promotion through Development → Testing → Production, with approval gates at each step.
- Integrations export as portable JSON templates — reusable, reviewable, auditable.
- **Samba never deletes a target record.** Not a gap — a stance. The highest-consequence failure mode in data integration shouldn't be one keystroke away.

It helps you run them

- Reads a failed execution and reports which stage broke, whether it's transient or permanent, and which mapping or HTTP call caused it.
- Analyzes deployment and cluster health — pods, database, Vault, scheduling, execution history, and a live Kubernetes view, captured in about five seconds.
- Every finding carries a certainty level — **confirmed, likely, possible, ruled out** — and cites the evidence behind it.

The AI cannot invent a diagnosis.

Findings come from a closed, versioned catalog of codes. The AI narrates the evidence — it does not get to make up what's wrong. On the authoring side, the same discipline: every generated plan is validated against a deterministic grammar before it reaches the runtime, and what ships is plain Python you read, test, and source-control. **Non-deterministic model, deterministic gate.** And it runs on your own Azure OpenAI endpoint, under your key, inside your tenant — no shared inference, ever.

RECOVERY YOU CAN HAND TO AN AUDITOR

- **Actions are signed, scoped, and reversible** — run pending migrations, re-seed configuration, rotate keys, re-register. Each with a defined blast radius, a rollback path, admin confirmation, and an audit-log entry.
- **Never free-form commands.** The system offers only the actions its current state warrants, and tells you what each one will **not** do before you confirm.
- **Diagnostic bundles carry no secrets or payloads** — schema metadata, counts, timestamps, and status flags only. Safe to attach to a ticket.