PROSPECT REVIEW | JUNE 2026

SABERRA

Platform Technical & Governance Documentation

Architecture Overview for Technical Reviewers, Data Privacy Officers, and Procurement Teams

Document Type: Prospect-Facing Architecture Overview

Version: June 2026 | Status: Draft for Human Review

saberra.com

1. What Saberra Is and Is Not

Saberra is an institutional memory system, not a student information system, a communication platform, or a data warehouse.

Saberra captures governance and operational content -- decisions made in meetings, tasks assigned in email threads, risks flagged in conversations, policies discussed in documents -- and structures that content into a human-reviewed, source-backed record your team can query and trust.

What Saberra Does Not Do

Saberra does not:

- Connect to student information systems (SIS), LMS platforms, or enrollment databases
- Access, store, or process student PII, grades, attendance, or behavioral records
- Integrate with DreamClass, PowerSchool, Canvas, or similar platforms unless explicitly scoped in writing
- Send, publish, or share captured content without human approval
- Train AI models on client data

Saberra's scope for education clients is limited to institutional and governance content: accreditation records, board or leadership decisions, operational policies, role assignments, risks, and strategic commitments -- the organizational memory of the institution, not the records of its students.

2. Platform Architecture

Saberra runs as three coordinated services within the client's dedicated cloud environment. Each service has a distinct responsibility and communicates only as required for its function.

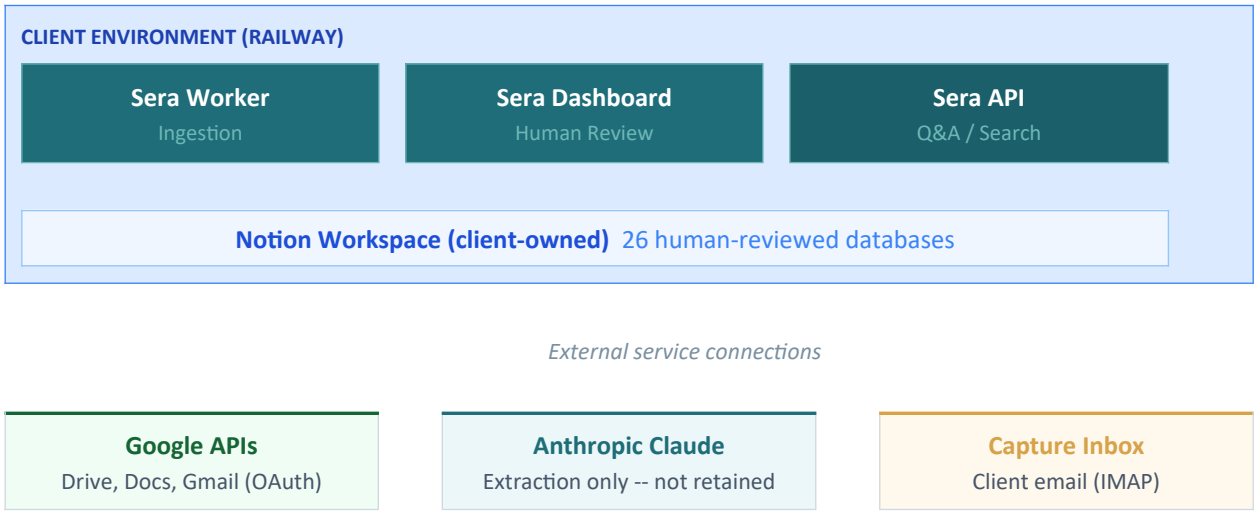

2a. Sera Worker -- Ingestion Service

Background process that polls the client's designated capture inbox every 3 minutes via IMAP. Classifies incoming emails, requests Google Drive access for meeting assets, and sends content to Claude for extraction. Writes only candidate/draft records -- never approved canon.

2b. Sera Dashboard -- Human Review UI

Password-protected web interface where authorized staff review, approve, reject, or modify AI-extracted candidate records. Nothing in Notion becomes trusted institutional memory until a human acts on it here.

2c. Sera API -- Q&A Service

HTTP service that answers natural-language questions exclusively from approved, human-reviewed Notion records. Cannot answer from unreviewed candidates. Cannot access external sources.

3. Information Flow

Every piece of content that enters Saberra follows this five-step path. No step can be bypassed by the system -- human review is structurally required before any record becomes trusted memory.

1

CAPTURE

Authorized sender forwards email or meeting link to the designated capture inbox
Sera Worker polls inbox via IMAP every 3 minutes
Messages not marked 'seen' until all processing completes (crash-safe design)

↓

2

CLASSIFICATION

Worker classifies message type: Google Meet recording/transcript, operational email, or system notification
Source Email record created in Notion as an audit trail entry
Routing determined: Meet path, Message path, or Access Request path

↓

3

AI-ASSISTED EXTRACTION (not AI decision)

Text content sent to Anthropic Claude API for structured extraction
Claude returns JSON candidates: decisions, tasks, risks, roles, policies, commitments, profiles, and more
ALL OUTPUTS ARE CANDIDATES ONLY -- nothing written as approved at this step
Sensitive content automatically flagged and routed to restricted Sensitive Review queue

↓

4

HUMAN REVIEW

Staff reviews candidates in Sera Dashboard
Approve: record becomes trusted institutional memory
Reject: record archived with rejection note
Modify: staff edits before approving
No AI-generated record becomes canon without this step

↓

5

RETRIEVAL

Approved records live in client-controlled Notion workspace
Sera API answers questions exclusively from approved records
Source citation included in every answer
If no reviewed record exists, Sera says so explicitly

> **Design guarantee:** The pipeline physically cannot write an approved record without a human action in the review step. This is not a policy -- it is a structural constraint in the architecture.

4. Data Residency and Ownership

All client data is stored in systems owned and controlled by the client. Saberra holds no copy of client data.

Data Type	Where It Lives	Who Controls Access
Captured emails / transcripts	Client Notion workspace	Client (Notion permissions)
Extracted candidate records	Client Notion workspace	Client
Approved institutional memory	Client Notion workspace	Client
Sensitive Review items	Separate admin-only Notion workspace	Designated admin only
Source email archive	Client Notion workspace	Client
Processing audit log	Client Notion workspace	Client
AI extraction input/output	Anthropic API (transient only)	Not retained by Anthropic *
Worker state / retry queue	Client Railway Postgres	Client Railway environment

** Anthropic's standard data processing agreement prohibits training on API customer data.*

> **Revoking access:** Removing the Notion integration token immediately terminates all Saberra write and read access. All records remain in the client's Notion workspace.

5. What Enters the System and How to Control It

Saberra is an opt-in capture system. Content enters only when someone explicitly sends it to the capture inbox or when a Google Meet asset notification is routed there.

Available Client Controls

Control	Description
Inbox access	Only staff with the capture email address can send content for ingestion. The address is not public.
Sensitive Review gate	Content flagged as sensitive (PII, financial data, legal matters) is automatically routed to the restricted Sensitive Review queue before any other processing.
Human review before canon	No extracted content becomes part of searchable institutional memory without explicit staff approval.
Granularity settings	Extraction depth (Essential / Standard / Full) is configurable per deployment.
Notion workspace permissions	Client controls who can view, edit, or access each Notion database independently of Saberra.

> Note for education clients: The capture inbox should only receive governance, accreditation, and operational content. Student-related content should never be forwarded to the capture address. This boundary requires clear staff training and is documented in the Selective Import Standard (Section 7).

6. Security and Privacy Guardrails

The following technical controls are implemented in the Saberra platform architecture. These are structural constraints, not operational policies.

Guardrail	Implementation
No student data access	No SIS/LMS integration. Capture inbox receives only what staff explicitly forward. No automated data collection.
AI cannot publish	All Claude outputs are candidates. Notion write access requires human approval. No direct AI-to-memory pathway exists.
Sensitive auto-routing	Regex and AI classification flags PII, financial data, and legal content for restricted review before any further processing.
Crash-safe ingestion	Messages not marked 'seen' in email until all processing completes. Notion dedup prevents reprocessing on retry.
Full audit trail	Every ingestion event written to Processing Events database with timestamp, source, and outcome. Immutable log.
Admin-only sensitive queue	Sensitive Review database lives in a separate Notion workspace with explicit permission grant required.
Transport security	All connections over TLS. No plaintext transmission of content at any point in the pipeline.
Credential isolation	All API keys and tokens stored as Railway environment variables. Never committed to source code or logs.
Client-controlled revocation	Removing Notion integration token immediately terminates all Saberra write and read access. Instant and complete.

7. Selective Import Standard

Saberra operates under a formal Selective Import Standard (adopted June 2026, version 1.0) that governs what content may be submitted for ingestion.

> **Core rule:** Only reviewed, current, and purposefully selected information may be submitted to Sera. Bulk imports of more than 5 documents require explicit approval from a designated Circle Lead or Data Steward.

Three Defined Roles

Role	Accountability
Data Steward	Accountable for what enters the system. Reviews and approves all bulk submissions before ingestion.
Bulk Import Approver	Circle Lead authority to approve large-scale ingestion (more than 5 documents at once).
Memory Auditor	Conducts periodic review of what is in the system for relevance, accuracy, and scope compliance.

This standard directly addresses the risk of inadvertently submitting sensitive, outdated, or out-of-scope content -- including content that might contain student information.

8. CCOS: Governance Architecture and Human Oversight

CCOS (Conscious Collective Operating System) is Saberra's internal governance framework, which Saberra helps clients implement and document. The governance architecture directly determines how data decisions are made, who is accountable for them, and how they are recorded and reviewable.

8a. What CCOS Draws From

Source	Contribution to CCOS
Holacracy	Roles over titles. Authority lives in defined role domains with explicit accountabilities and decision rights, not in organizational hierarchy.
Sociocracy (S3)	Consent-based decisions. Double-linking between circles so information flows both ways. Decisions require consent from those affected.
Cooperative models	Distributed ownership. Transparency as a structural commitment. No single stakeholder can unilaterally close the system to others.

8b. How CCOS Manifests in Saberra's Design

Governance Principle	How It Is Implemented in Saberra
AI proposes; humans decide	All AI outputs are candidates. No record becomes canon without explicit human approval in the review step.
Defined role accountability	Every circle has a Circle Lead and Rep Steward. Data decisions are made by defined roles, not by the AI or the vendor.
Consent before canon	Canon Change Requests require human consent before policy or governance records are ratified.
Transparency Ledger	All formal governance actions recorded in the CCOS Ledger with full status history: Draft, Pending Review, Ratified.
Sensitive Review gate	Restricted information requires explicit authority to access -- mirrors the principle that access follows role, not curiosity.
Double-linking	Data concerns can flow from any circle upward through the Rep Steward role for governance review.
Distributed authority	No single person -- including the founder -- can approve canon records. The system prevents that concentration.

8c. The Governance Decision Flow

The following flow applies to data policy changes, role definition changes, sensitive review handling procedures, and any governance decision that affects how information is managed.

1. Tension identified -- by any role holder in any circle

↓

2. Proposal drafted -- by the role closest to the tension

↓

3. Objection round -- any role may object if the proposal causes harm, confusion, or trust risk

↓

4. Consent reached -- 'safe enough to try until next review'

↓

5. Decision recorded in Transparency Ledger by Sera as a Draft candidate

↓

6. Human review and ratification -- a designated reviewer approves the record

↓

7. Canon -- reviewable, citable, source-backed institutional memory

9. For Education Clients: Scope Boundary

Saberra's role in an education deployment is strictly defined. The following clarifies what is and is not within scope for the standard Saberra deployment.

In Scope

- ✓ Accreditation governance records
- ✓ Board and leadership decisions
- ✓ Operational policies and procedures
- ✓ Role assignments and accountability structures
- ✓ Strategic risks and institutional commitments
- ✓ Institutional knowledge (not student knowledge)

Out of Scope

- ✗ Student records of any kind
- ✗ Parent or guardian communications
- ✗ Behavioral or disciplinary records
- ✗ Health or wellness data
- ✗ Financial aid or enrollment data
- ✗ Content governed by FERPA, COPPA, or equivalent

Student data governance

Attendance, grades, demographics, and enrollment data are managed by the institution's student information system (e.g. DreamClass) which operates independently from Saberra. There is no integration between Saberra and the SIS unless separately and explicitly scoped in a written agreement.

10. Summary: The Trust Boundary

The following is Saberra's core trust commitment, as stated in its public documentation. It is enforced structurally -- not by policy alone.

> ***"Client records live in client-controlled systems. AI creates candidates. Humans review them. Sera answers only from reviewed records."***

For a prospect or data privacy officer concerned about student data, the relevant assurances are:

Core Trust Commitments

1. Saberra does not touch student data unless explicitly scoped and agreed in writing
2. Everything that enters the system requires explicit staff action to get there
3. Nothing AI-extracted becomes trusted memory without human approval
4. Sensitive content is automatically gated behind a restricted review queue
5. The client owns and controls all records in their Notion workspace
6. Access can be revoked instantly by removing the integration token

> **This document is a draft for human review.** It does not constitute a legal data processing agreement or privacy certification. A formal Data Processing Agreement (DPA) is available upon request at saberra.com.