Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

PREE AI: A Zero-Trust, Institution-Sovereign Architecture for Multilingual and Privacy-Preserving Educational Intelligence

Integrating Federated Learner Analytics, Curriculum-Grounded Retrieval, Multilingual Intent Understanding, and Verifiable Academic Credentials

Aditya Mandloi, Ashok Kumar Gupta, and Govind Maheshwari, Md. Kaifee
School of Data Science and Forecasting, Devi Ahilya Vishwavidyalaya, Indore, India

Abstract **- Educational AI systems increasingly combine conversational tutoring, learning analytics, and digital academic services. Yet these capabilities are commonly designed as separate services, creating a fragmented trust model: student records may be centralized for analytics, generative answers may be weakly grounded in prescribed curricula, multilingual and code-mixed questions may be inconsistently interpreted, and academic eligibility checks may require disclosure of personal records. This paper proposes PREE AI (Predictive & Responsive Education Ecosystem), a zero-trust, institution-sovereign architecture that coordinates four complementary capabilities: (i) cross-silo federated learner-risk analytics with a differentially private training option; (ii) curriculum-constrained hybrid retrieval combining BM25 lexical ranking and FAISS dense-vector search through Reciprocal Rank Fusion; (iii) multilingual and code-mixed intent routing using an IndicBERT-based classifier; and (iv) Groth16 zero-knowledge proofs for selective verification of academic eligibility. A translation and response-caching layer supports multilingual interaction, while an API gateway coordinates policy enforcement and service communication. The paper's central contribution is architectural: it defines trust boundaries, data flows, subsystem interfaces, and a feasible coordination model for delivering learner support without requiring a single centralized repository of raw student records. We provide a formal system model and requirement set, interface contracts, workflow algorithms, an information-flow and threat analysis, governance safeguards, and a metrics-and-protocol specification for future validation. The framework is presented as a technically motivated architecture; population-level learning gains and subsystem performance figures are not claimed without a reproducible implementation and evaluation.**

Keywords—Artificial Intelligence in Education; Federated Learning; Differential Privacy; Hybrid Retrieval-Augmented Generation; BM25; FAISS; IndicBERT; Zero-Knowledge Proofs; Groth16; Educational Data Governance.

I. INTRODUCTION

Digital learning environments now support tutoring, formative assessment, content discovery, and student-success monitoring [1]. However, integrating these functions into a coherent educational service is difficult, particularly in linguistically diverse settings and across institutions that operate independent information systems. In India, students may formulate questions in English, an Indic language, or code-mixed forms such as Hinglish. A tutoring system that assumes English-only interaction can introduce friction, while a translation-only wrapper may fail to preserve curriculum-specific terminology, equations, and local educational context [2]–[4].

A second challenge concerns the use of learner data for early support. Academic performance, engagement patterns, attendance, and assignment behavior can help identify students who may benefit from timely intervention [5], [6]. Yet a centralized analytics repository creates an attractive concentration of sensitive records and can be difficult to reconcile with data-minimization principles and the different legal obligations that apply across jurisdictions [7]–[10]. Institutions may also be unwilling or unable to share identifiable learner-level data, even when collaboration could improve the robustness of risk models.

A third challenge arises from generative tutoring. A fluent answer is not necessarily a curriculum-correct answer. Dense semantic retrieval can identify conceptually related passages, but it may underweight exact identifiers, symbols, formulae, named theorems, and spelling variants. Conversely, lexical retrieval alone may fail when a learner paraphrases a concept or asks in a code-mixed form. A further challenge appears in inter-institutional academic workflows: a receiving institution or scholarship administrator may need to verify whether a learner satisfies a condition without receiving the learner's complete academic record.

These issues are often addressed independently. PREE AI is proposed as a unified architecture in which learner-risk modeling, curriculum-grounded tutoring, multilingual intent understanding, and credential verification operate behind explicit institutional trust boundaries. The objective is not to claim a new federated-learning algorithm, retrieval algorithm, language model, or proof system. Instead, the research contribution is a coordinated design that specifies how these established techniques can interoperate in an educational setting, what information crosses each boundary, and which responsibilities remain local.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

A. Research Question and Scope

The guiding research question is: How can an educational platform coordinate multilingual tutoring, collaborative learner-risk analytics, curriculum-grounded generation, and privacy-preserving academic verification while preserving institutional control over student records and curriculum resources? PREE AI addresses this question at the architectural and protocol-design level. It is intended as a reference architecture and implementation blueprint, not as evidence that the proposed system has already produced a measured improvement in retention or academic achievement.

B. Contributions

C1. An institution-sovereign learning-analytics design in which participating institutions retain raw learner records locally and exchange model updates under an explicitly defined federation policy.

C2. A curriculum-constrained hybrid retrieval design that combines BM25 lexical evidence with FAISS (Facebook AI Similarity Search) dense retrieval and fuses their rankings using Reciprocal Rank Fusion (RRF).

C3. A multilingual interaction pipeline that uses IndicBERT-based intent classification to route student requests to tutoring, assessment, support, or administrative workflows.

C4. A selective academic-eligibility verification design using Groth16 zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs), with public policy inputs separated from private student attributes.

C5. A cross-layer trust model and end-to-end workflow that clarify service boundaries, data minimization, failure handling, and the limits of the proposed privacy guarantees.

C6. A requirement-traceable specification of interface contracts, per-party information flows, governance safeguards, and a reproducible evaluation protocol against which future implementations can be audited.

II. RELATED WORK AND RESEARCH GAP

Research in AI-supported education includes intelligent tutoring systems, knowledge tracing, educational data mining, conversational agents, multilingual language technologies, and privacy-preserving computation. Classical systems such as AutoTutor and Cognitive Tutor established important foundations for adaptive instructional dialogue and domain-specific support [11], [12]. Knowledge-tracing research models the evolution of learner mastery over time [13]. Educational data-mining and academic-analytics studies have demonstrated the value of behavioral and assessment data for understanding student progress [6], [14], [15].

More recent conversational systems use large language models (LLMs) to support open-ended interaction. Their usefulness in education depends not only on fluency but also on age-appropriate behavior, curriculum alignment, source traceability, and appropriate escalation when a question cannot be answered reliably [16]. Retrieval-augmented generation (RAG) offers a practical way to condition generation on selected source material [17]. Dense retrieval supports semantic matching, while lexical methods such as BM25 are useful for exact terms and identifiers [18]. Combining these retrieval signals can reduce dependence on a single representation of relevance [19].

Multilingual language models and translation systems, including NLLB, IndicTrans2, and IndicBERT-related resources, provide building blocks for language coverage and cross-lingual understanding [2]–[4]. Federated learning offers a way to train shared models without collecting all training examples in one location [20], while differential privacy can bound information leakage under a specified adjacency definition and accounting procedure [21], [22]. Zero-knowledge proof systems such as Groth16 can establish that a statement about private witness data satisfies a public relation without revealing the witness itself [23].

The research gap addressed here is therefore not the absence of any one of these technologies. It is the lack of a clearly specified educational architecture that coordinates (a) institution-local learner analytics, (b) institution- or board-specific curriculum grounding, (c) multilingual query interpretation, and (d) selective credential verification under a common trust model. PREE AI makes this coordination explicit and identifies the technical and governance assumptions needed for implementation. Table I contrasts common design patterns with the PREE AI response.

Dimension	Common design pattern	PREE AI architectural response
Learner analytics	Centralized institutional or platform-level data aggregation	Local data custody with cross-silo model coordination
Tutoring knowledge	Single retrieval channel or model-only generation	BM25 + FAISS retrieval with rank fusion and source metadata
Language interaction	Language detection or translation treated as an outer wrapper	Intent-aware routing for multilingual and code-mixed requests
Academic verification	Disclosure or exchange of complete records	Proof of a narrowly defined eligibility predicate
System governance	Security controls distributed across services	Explicit trust boundaries, policy gateway, and local institutional control

TABLE I
COMPARISON OF COMMON DESIGN PATTERNS AND THE PREE AI ARCHITECTURAL RESPONSE

III. PROBLEM FORMULATION AND DESIGN REQUIREMENTS

This section fixes the system model, notation, requirements, and risk classes against which the architecture is specified. The formulation is qualitative by design: it states what each component must guarantee and what it is explicitly not asked to guarantee, so that a later implementation and its evaluation can be checked against stated claims rather than against an aggregate score.

A. System Model and Notation

We consider a federation of K institutions. Institution k operates a learner store D_k, a curriculum corpus C_k, an identity

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

provider, and optionally a credential-issuing authority. Learners reach PREE AI through the gateway of their home institution. Three external roles interact with the federation: a coordinator that aggregates model updates, relying institutions that request eligibility proofs, and model providers that host translation or generation models when these are not run locally.

The data in the federation fall into four classes with different disclosure rules: raw learner records, local model updates, curriculum artifacts, and proof artifacts. Only the last three may cross an institutional boundary, and each does so in a restricted form. Table II fixes the symbols used in the remainder of the paper.

Symbol	Meaning
K, k	Number of institutions; institution index
D_k, n_k	Local learner dataset of institution k and its size
C_k	Versioned curriculum corpus approved at institution k
θ_t	Global learner-risk model parameters at federated round t
C, σ	L2 clipping bound and Gaussian noise multiplier
q, T_i	Learner query and the i-th indexed curriculum passage
x, w	Public inputs and private witness of a credential circuit
π	Groth16 proof (π_A, π_B, π_C)
u, a, d	Requesting user, requested action, and data class

TABLE II
NOTATION USED THROUGHOUT THE PAPER

B. Design Requirements

Seven requirements, R1–R7, are derived from the four challenges in the Introduction and from constraints that institutions commonly impose on shared infrastructure. Table III states each requirement and the mechanism through which the architecture is intended to satisfy it. Every requirement is testable in principle: it can be associated with an inspection, a measurement, or a failure-injection experiment in the evaluation protocol described later.

ID	Requirement	Intended mechanism
R1	Raw learner records never leave the institution	Local BEACON training; gateway egress filter
R2	Every cross-boundary message follows a declared schema	Interface contracts; payload validation
R3	Answers are traceable to approved, versioned curriculum sources	Hybrid retrieval with source metadata; citations
R4	English, Indic, and code-mixed requests are interpreted consistently	IndicBERT router; translation cache
R5	Eligibility is verifiable without disclosing the record	Groth16 predicate proofs
R6	Decisions are auditable without duplicating learner content	Audit layer with content-free logs
R7	Failure of one subsystem neither leaks data nor blocks others	Separated workflows; abstain, clarify, deny-by-default

TABLE III
DESIGN REQUIREMENTS AND INTENDED MECHANISMS

C. Trust Assumptions and Risk Classes

The architecture does not assume that all parties are equally trustworthy. Five risk classes are considered:

- **A1, honest-but-curious coordinator:** follows the federation protocol but may inspect every update it receives.
- **A2, network observer:** may observe or alter traffic between institutions, relying parties, and model providers.
- **A3, curious relying institution:** may seek to learn more than the eligibility outcome from proofs and public inputs.
- **A4, learner-side misuse:** a learner's queries, uploads, or presented credentials may, intentionally or inadvertently, contain injected instructions or be a replayed or invalid credential.
- **A5, compromised or misconfigured internal service:** holds legitimate network access but exceeds its intended privileges.

A fully compromised credential issuer, physical compromise of institutional key storage, and side channels on proving hardware are outside the scope of the guarantees discussed here; their consequences are stated as residual risks in the security analysis.

IV. PREE AI ARCHITECTURE

PREE AI is organized as a layered, modular architecture (Fig. 1). Each subsystem can be deployed independently, but all services follow a shared policy model enforced at the institutional gateway. The architecture is designed to support multiple institutions, each of which may maintain its own learner database, curriculum corpus, identity provider, and credential-issuing authority.

A. Architectural Layers

1) *Interaction layer:* student, teacher, counselor, and administrator interfaces. It accepts text or supported input modalities and presents tutoring responses, risk alerts, and verification outcomes.
2) *Policy and application programming interface (API) gateway:* authenticates requests, applies role- and purpose-based authorization, validates payloads, rate-limits calls, removes or tokenizes unnecessary identifiers, and routes requests to permitted services.
3) *Intelligence layer:* multilingual translation/cache service based on neural machine translation (NMT), IndicBERT intent classifier, hybrid retrieval service, federated BEACON coordinator, and credential proof/verifier services.
4) *Institutional data layer:* locally controlled learner records, assessment and engagement data, curriculum documents, model checkpoints, and credential-issuer records. Data are not assumed to be pooled centrally.
5) *Audit and governance layer:* records policy decisions, model versions, proof-verification outcomes, source references, consent or lawful-basis metadata where applicable, and administrative actions without unnecessarily duplicating learner content.

The gateway is a coordination point, not a universal data lake. A tutoring request should reach only the curriculum retrieval and language services required to answer it. A risk-model update should be computed within the institution's environment. A

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

credential request should disclose only the public statement and proof needed for verification.

B. Trust and Data-Residency Principles

PREE AI adopts four design principles. First, raw learner-level records remain under the control of the institution that collected them. Second, shared model updates are treated as sensitive messages rather than inherently anonymous information. Third, curriculum content is indexed and filtered according to the institution's permitted syllabus and source versions. Fourth, academic verification is predicate-specific: a verifier learns whether a defined condition holds, not the complete student record. These principles guide the interfaces described below.

Two design consequences follow. Policy is enforced at the institutional gateway rather than inside individual subsystems, so a change in consent, retention, or disclosure rules is applied uniformly and recorded once in the audit layer. Interfaces also carry only the fields that a subsystem needs for its own function, which keeps a fault or compromise in one component from exposing data that belong to another.

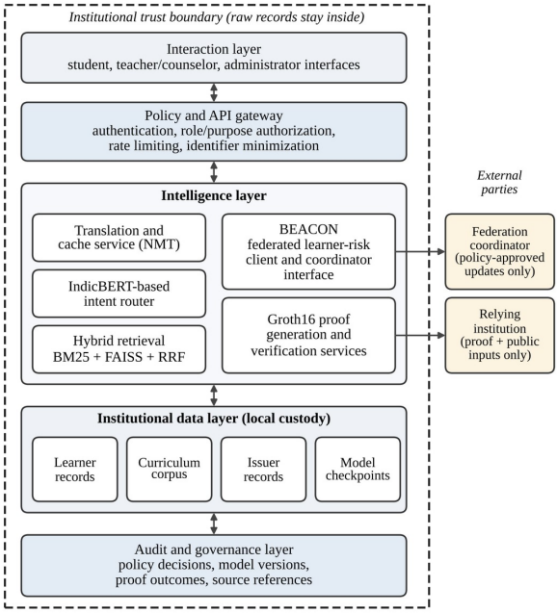


Fig. 1. Layered PREE AI architecture. Raw learner records, curriculum corpora, issuer records, and model checkpoints stay inside the institutional trust boundary; only policy-approved model updates and proofs with public inputs cross it.

C. Core Subsystem Coordination

Table IV summarizes the responsibilities that remain local to each institution and the information that may cross a trust boundary.

Subsystem	Local responsibility	Cross-boundary exchange
BEACON	Feature preparation, local training, clipping/noise, local evaluation	Policy-approved model updates and aggregate diagnostics
Hybrid retrieval	Curriculum ingestion, chunking, indexing, access filtering	Query and retrieved excerpts only when authorized
IndicBERT/NMT	Intent prediction, language handling, translation cache	Intent label, confidence, and minimum necessary text
Groth16	Credential issuance, witness construction, proof generation	Proof and explicitly defined public inputs

TABLE IV
LOCAL RESPONSIBILITY AND CROSS-BOUNDARY EXCHANGE BY SUBSYSTEM

D. Deployment Topologies

The same services support three deployment patterns (Fig. 2). In a *single-institution* deployment, federation is disabled and BEACON trains and serves models locally; this is the lowest-risk starting point and exercises the gateway, retrieval, routing, and credential services without any cross-boundary model exchange. In a *consortium federation*, several institutions train a shared risk model through a neutral coordinator, and the trust assumptions and security analysis apply in full. In a *board-hosted curriculum* pattern, an examination board or education authority publishes signed, versioned curriculum bundles that institutions ingest into their local indexes; learner data remain local, but the board becomes a trusted source of curriculum content and its signing keys become a governance concern.

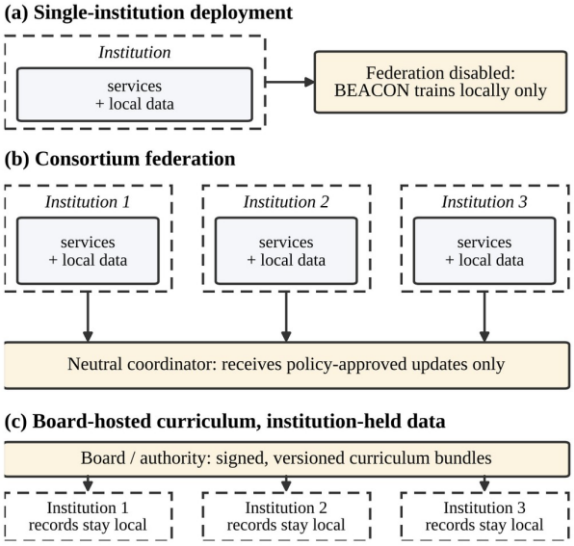


Fig. 2. Deployment patterns supported by the architecture: (a) single institution with federation disabled, (b) consortium federation through a neutral coordinator, and (c) board-hosted curriculum bundles with institution-held learner data.

These patterns are not mutually exclusive. An institution may begin with (a), join (b) for risk analytics once feature definitions are harmonized, and adopt (c) for curriculum distribution independently of whether it participates in federated training. The gateway enforces the same authorization policy in all three cases.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

V. CORE ARCHITECTURAL COMPONENTS

A. BEACON: Institution-Sovereign Federated Learner-Risk Analytics

BEACON is the proposed learner-risk component. Its purpose is to allow institutions to improve a shared risk model without exporting raw student telemetry. At institution k, a local feature pipeline derives a learner representation from permitted data. The conceptual feature groups include academic progress A, engagement E, stress-related indicators S where ethically and lawfully collected, and submission punctuality B. Each feature is normalized to a documented range and accompanied by provenance and missingness information.

A simple interpretable risk score can be expressed as:

$$R_k(u) = \sigma(\lambda_A(1 - A_u) + \lambda_E(1 - E_u) + \lambda_S S_u + \lambda_B(1 - B_u)) \quad (1)$$

where u denotes a learner, σ is a logistic or other bounded mapping, and λ_A, λ_E, λ_S, and λ_B are learned or institutionally calibrated weights. This expression is a design-level scoring formulation; feature definitions, labels, and calibration must be validated for each deployment. In particular, psychological or stress-related indicators should not be inferred from intrusive surveillance or treated as clinical diagnoses. They should be excluded unless collection is justified, proportionate, consented or otherwise lawfully grounded, and reviewed by appropriate institutional authorities.

For federated training, institution k minimizes a local objective $F_k(\theta)$ over its own dataset D_k. A conventional weighted federated averaging (FedAvg) update [20] is:

$$\theta_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} \theta_{t+1}^{(k)}, \quad n = \sum_{k=1}^{K} n_k \quad (2)$$

where K is the number of participating institutions, n_k is the number of local training examples at institution k, and $\theta^{(k)}_{t+1}$ is the locally updated model. A privacy-enhanced implementation may clip per-example or per-user gradients to a declared L2 bound C and add calibrated Gaussian noise before release [22]. However, a numerical privacy guarantee cannot be inferred from the presence of Gaussian noise alone. The implementation must specify the adjacency unit (for example, learner-level or record-level), clipping location, sampling scheme, noise multiplier, number of rounds, participation assumptions, and a valid privacy accountant. The value of ε is therefore a deployment parameter to be calculated and reported after the training protocol is fixed; this paper does not assert a verified ε value.

BEACON also requires safeguards against non-independent and identically distributed (non-IID) institutional data, a recognized open problem in federated learning [36]. Institutions may differ in grading scales, course structures, student demographics, attendance policies, and definitions of dropout. Before aggregation, feature semantics and label definitions must be harmonized or explicitly modeled. Local validation should be retained, and global model quality should be reported alongside institution-wise performance to avoid hiding poor outcomes for smaller or underrepresented institutions.

B. BEACON Feature Schema and Label Definition

Federated training is only meaningful if participating institutions agree on what is being predicted and from what. Table V summarizes the feature groups introduced above, the kind of local fields that may populate them, and the harmonization rule applied before training. The outcome label (for example, non-completion within a fixed observation window) must be defined identically across institutions, and every feature must be computed only from data available at prediction time. Otherwise temporal leakage can inflate apparent accuracy and hide the cases the system is meant to catch.

Missing values are modeled explicitly rather than imputed silently, since the pattern of missingness can differ systematically between institutions and carries its own privacy implications. Each institution reports the proportion of missing values per feature group in its local diagnostics so that the coordinator can detect schema drift without seeing any record.

C. Curriculum-Constrained Hybrid Retrieval

The PREE AI retrieval subsystem is designed to ground tutoring responses in institution-approved curriculum material (Fig. 3). Documents are versioned, access-controlled, and divided into passages with metadata such as board, subject, grade, chapter, source, page, and effective date. A query is filtered by the learner's authorized curriculum context before retrieval.

Passages should follow natural curriculum boundaries such as chapters, sections, and worked examples, so that each one keeps a stable identifier tied to its source, page, and edition. A passage that cannot be traced to an approved source is not eligible for retrieval, and a superseded edition is withdrawn from the index while remaining available to the audit layer for reconstructing earlier answers.

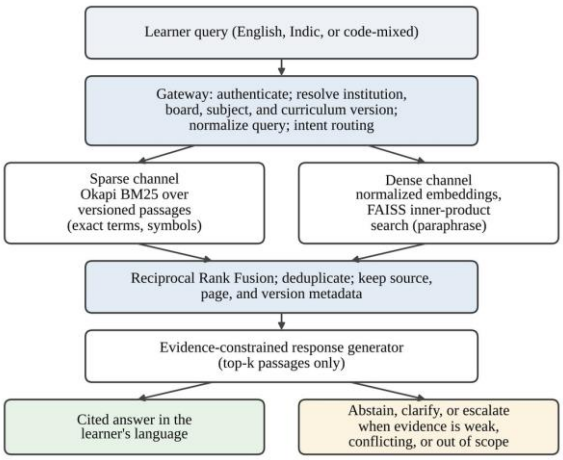


Fig. 3. Hybrid retrieval and grounded-response workflow with an explicit abstain, clarify, or escalate path.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Group	Example local fields	Harmonization rule
Academic (*A*)	Assessment scores, grades, credit load	Map local grading scheme to a documented range; retain scheme identifier
Engagement (*E*)	Activity counts, resource interactions	Normalize by course exposure time; record missingness
Punctuality (*B*)	Submission time relative to deadline	Use relative lateness; align deadline definitions
Stress (*S*)	Optional consented wellbeing instrument	Disabled by default; excluded unless justified and reviewed
Label	Non-completion in a fixed window	Shared definition and window agreed before training

TABLE V
BEACON FEATURE GROUPS AND HARMONIZATION RULES

The sparse channel uses Okapi BM25 [18] to identify exact lexical matches. For query terms q_j and passage T_i, a conventional BM25 score is:

$$\mathrm{BM25}(q, T_i) = \sum_j \mathrm{IDF}(q_j) \frac{f_{ij}\,(k_1 + 1)}{f_{ij} + k_1\,(1 - b + b\,L_i)} \tag{3}$$

where IDF is the inverse document frequency, f_{ij} is the frequency of q_j in T_i, $L_i = |T_i|/\mathrm{avgdl}$ is the length of T_i relative to the average passage length avgdl, and k_1 and b are tuning constants. The dense channel encodes the query and passages into normalized vectors and searches a FAISS index [24] using inner-product similarity, which is equivalent to cosine similarity for normalized vectors:

$$\mathrm{sim}(q, T_i) = \mathbf{q}^{\top}\mathbf{v}_i, \quad \|\mathbf{q}\|_2 = \|\mathbf{v}_i\|_2 = 1 \tag{4}$$

Because BM25 and dense similarity scores are not naturally calibrated to the same scale, PREE AI combines their ranked lists using Reciprocal Rank Fusion [19]:

$$\mathrm{RRF}(T_i) = \frac{1}{\kappa + r_{\mathrm{BM25}}(T_i)} + \frac{1}{\kappa + r_{\mathrm{FAISS}}(T_i)} \tag{5}$$

where r denotes a rank and κ is a smoothing constant selected during validation. The fused ranking supplies the top-k passages to the response generator. Each passage should retain source and page metadata so that the system can display citations or trace the answer to the approved curriculum.

The novelty claim is not that BM25, FAISS, or RRF is new. It is the use of a dual-evidence retrieval design as a curriculum constraint within an institution-sovereign educational workflow. Exact formulae, theorem identifiers, chemical notation, and named concepts can benefit from lexical matching, while paraphrased or code-mixed questions can benefit from semantic retrieval. The system should abstain or ask for clarification when the retrieved evidence is weak, conflicting, outdated, or outside the learner's authorized syllabus.

D. IndicBERT-Based Multilingual Intent Routing

The interaction layer uses an IndicBERT-based classifier [4] to estimate the intent of a learner request. Candidate intents may include concept explanation, worked example, practice question, answer checking, curriculum lookup, administrative query, wellbeing support, and escalation to a human educator. The classifier produces a probability distribution over intent classes:

$$P(I \mid x) = \mathrm{softmax}(W\mathbf{h}_{[\mathrm{CLS}]} + \mathbf{c}) \tag{6}$$

where x is the input utterance, $\mathbf{h}_{[\mathrm{CLS}]}$ is the contextual representation produced by the encoder, and W and $\mathbf{c}$ are learned parameters. The predicted intent is used for routing, not as a substitute for domain reasoning. Confidence thresholds and an out-of-scope path [25] should be included so that ambiguous requests are not forced into an incorrect category.

Code-mixed examples are important because learners may combine an Indic language with English technical vocabulary, mathematical notation, or transliteration. Training and evaluation data should represent language pairs, dialect variation, spelling variation, and educational subject areas. The architecture permits an NMT service or translation cache to support response rendering, but translation quality and latency must be evaluated separately from intent-classification quality. A cache should be keyed carefully to avoid exposing sensitive query content or reusing a response across users with different permissions or curriculum contexts.

E. Groth16-Based Verifiable Academic Credentials

The credential subsystem supports a narrow class of cross-institutional checks. For example, a student may need to prove that a grade point average (GPA) meets a published threshold without disclosing the exact GPA or full transcript. A credential issuer signs or otherwise authenticates the underlying academic claim. A prover constructs a Groth16 proof $\pi = (\pi_A, \pi_B, \pi_C)$ that a circuit relation is satisfied by private witness values and public policy inputs [23].

Let x contain public inputs such as a threshold, issuer or institution identifier, credential schema version, and expiry condition. Let w contain private values such as the student's GPA, credential data, and a blinding value. The circuit checks issuer authenticity and the predicate GPA ≥ threshold, together with any required validity constraints. The proof π is verified against a verification key and the declared public inputs. A schematic Groth16 verification relation is:

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(\mathrm{vk}_x, \gamma) \cdot e(\pi_C, \delta) \tag{7}$$

where e is a bilinear pairing, α, β, γ, and δ are verification-key elements, and vk_x is the verification-key linear combination derived from public inputs. The exact circuit, field representation, credential binding, and pairing equation must follow the selected Groth16 implementation and library. The verifier should return only the result and the public policy context required by the relying party.

Zero knowledge does not by itself guarantee that the complete workflow reveals nothing. Public inputs, issuer identifiers, timestamps, repeated presentations, network metadata, and compromised issuer keys can still expose information. The design therefore requires proof freshness, replay protection, key management, issuer authentication, revocation or expiry handling, and careful minimization of public inputs. Groth16 also requires a trusted setup for the circuit; multi-party setup ceremonies can reduce reliance on any single participant [37], and setup

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

governance and verification-key rotation remain operational considerations.

F. Credential Circuit Specification

To make the credential design auditable, the statement proved by the circuit is specified separately from the proof system. Table VI lists the public inputs, private witness, and constraints for the GPA-threshold example. The grade-point average is encoded as a fixed-point integer so that the comparison with the threshold is a bounded integer range check, and a verifier-supplied nonce is bound into the public inputs so that a proof generated for one request cannot be replayed for another.

Changing any constraint changes the verification key and therefore requires a new setup; circuit versions should be recorded in the audit log and in the credential schema version so that verifiers can reject proofs produced for superseded circuits.

G. Translation, Caching, and Orchestration

The translation layer supports language detection, optional translation, response localization, and cache reuse. It should preserve technical tokens, formulae, code fragments, names, and curriculum terminology. Cache entries should include model version, language pair, curriculum context, and policy scope. Sensitive learner text should not be retained by default; where caching is permitted, retention and access controls must be explicit.

The gateway orchestrates the sequence: authenticate the learner, classify intent, determine the permitted workflow, retrieve curriculum evidence when needed, invoke the response model, apply output checks, and return a localized answer with source metadata. Risk analytics and credential verification are separate workflows and should not be invoked merely because a learner asks an ordinary tutoring question.

H. Curriculum Index Lifecycle

Retrieval quality depends on index hygiene as much as on ranking. Each source document is ingested with its board, subject, grade, edition, and effective date; passages inherit this metadata and a content hash. When a syllabus changes, the superseded edition is marked inactive rather than silently overwritten, so that earlier answers remain traceable in the audit log. Changing the embedding model invalidates the dense index and requires re-embedding all passages, while the lexical index is unaffected; the two indexes are therefore versioned independently and queried only when their versions are mutually compatible. Deletion requests for source material propagate to both indexes and to any cached responses derived from them.

Element	Content
Public inputs *x*	Threshold; issuer identifier; credential schema version; expiry epoch; verifier nonce
Private witness *w*	Fixed-point GPA; issuer signature on the credential; blinding value; learner secret
Constraints	(i) issuer signature verifies over the committed credential; (ii) GPA ≥ threshold as a range check; (iii) credential unexpired at the stated epoch; (iv) nonce bound to the proof
Output	Proof π and the public inputs; no other attribute of the credential

TABLE VI
CIRCUIT SPECIFICATION FOR THE GPA-THRESHOLD PREDICATE

I. Interface Contracts and Policy Evaluation

Requirement R2 is enforced through explicit interface contracts. Table VII lists, for each cross-component interface, the fields a message may carry and the content it must never carry. The gateway validates payloads against these schemas and rejects messages that include undeclared fields. Authorization is evaluated per request as the conjunction

$$\mathrm{Allow}(u, a, d) = \mathbf{1}[R(u, a) \wedge P(a, d) \wedge S(u, d)] \quad (8)$$

where $R(u,a)$ holds if the role of user u permits action a, $P(a,d)$ holds if the declared purpose of a is permitted for data class d, and $S(u,d)$ holds if d lies within the institutional and curriculum scope of u. A request is denied by default when any term cannot be evaluated.

Interface	May carry	Must never carry
Gateway → retrieval	Normalized query; board, subject, grade, version filters	Learner identifier; unrelated profile data
Gateway → intent router	Query text; language tag	Risk scores; credential data
BEACON client → coordinator	Clipped (and, if configured, noised) update; schema version	Raw records; per-learner gradients
Prover → verifier	Proof π; declared public inputs *x*	Witness *w*; full transcript
Services → audit log	Decision, versions, source IDs, timestamps	Learner text, by default

TABLE VII
INTERFACE CONTRACTS FOR CROSS-COMPONENT MESSAGES

VI. END-TO-END ALGORITHMS AND WORKFLOW

A. Federated BEACON Training Workflow

1) Coordinator publishes a versioned model, feature schema, training configuration, and participation policy.
2) Each participating institution validates the schema and trains on its locally held learner records.
3) Local updates are clipped and, when the selected privacy protocol requires it, perturbed using a calibrated mechanism.
4) Only approved updates and non-identifying diagnostics are transmitted through an authenticated channel.
5) The coordinator aggregates eligible updates and distributes the next model version.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

6) Each institution evaluates the received model locally and may reject or roll back a model that fails its safety or performance criteria.

Fig. 4 illustrates one federated round and the points at which data cross the institutional boundary: only the clipped, optionally noised update leaves the institution, and only a versioned model returns.

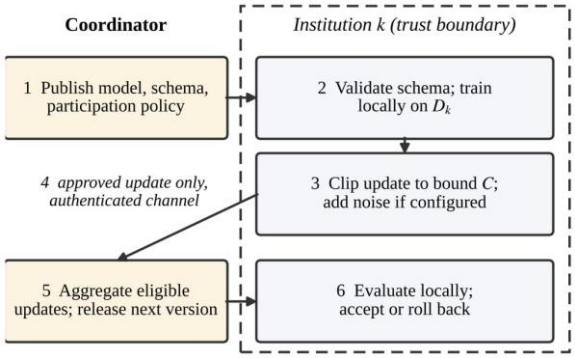


Fig. 4. One federated BEACON round. Steps are numbered as in the workflow list; only step 4 crosses the institutional trust boundary outbound.

B. Hybrid Retrieval and Grounded Response Workflow

1) Authenticate the learner and resolve the permitted institution, board, subject, and curriculum version.
2) Normalize the query while preserving mathematical symbols, technical terms, and relevant transliterations.
3) Retrieve candidate passages independently using BM25 and FAISS under the same access-control filters.
4) Fuse the ranked candidates using RRF, remove duplicates, and retain source, page, and version metadata.
5) Pass the selected evidence to the response generator with instructions to answer only to the extent supported by the evidence.
6) Return the answer with citations; abstain, clarify, or escalate when evidence is insufficient or contradictory.

C. Credential Verification Workflow

1) The relying institution publishes the eligibility predicate and the required public inputs.
2) The credential issuer authenticates the academic record and issues a verifiable credential to the learner.
3) The learner's device or authorized prover constructs a proof for the requested predicate without sending the private witness.
4) The verifier checks the proof, public inputs, issuer trust, freshness, and revocation status.
5) The relying institution receives a minimal valid/invalid outcome and the policy context, not the complete transcript.

These workflows are intentionally separated. A shared gateway can enforce common authentication and audit policies, but the data exchanged by each workflow differs. This separation reduces unnecessary coupling and limits the consequences of a compromised service.

Fig. 5 shows the message sequence for a single tutoring request. The audit log receives only the policy decision, model and index versions, and source identifiers, and the learner-facing answer is always accompanied by its sources or by an explicit abstention.

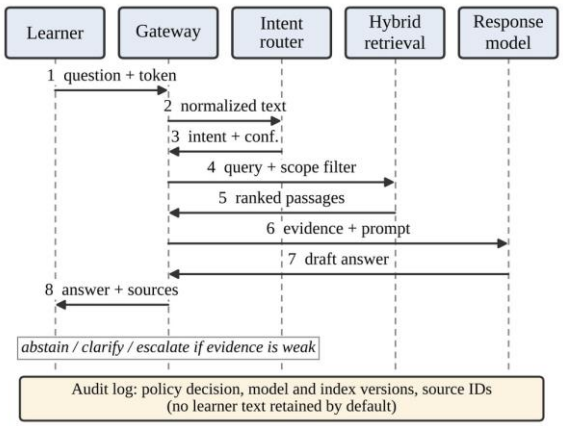


Fig. 5. Message sequence for a curriculum-grounded tutoring request. Numbered arrows are in-order calls; the gateway is the only component that talks to both the learner and the intelligence services.

D. Failure Handling and Graceful Degradation

Requirement R7 is addressed by specifying a safe default for each failure mode rather than by assuming that all services remain available:

- **Translation unavailable:** respond in the language of the source passage with a visible notice; do not fall back to an external service that was not authorized for the request.
- **Weak or conflicting evidence:** abstain, ask a clarifying question, or escalate to a human educator rather than generating an unsupported answer.
- **Coordinator unreachable:** continue serving the last accepted risk model and queue local updates; no raw data are sent as a workaround.
- **Prover or verifier failure:** deny by default and report an indeterminate outcome, never a positive one.
- **Low intent confidence:** route to a clarification prompt or the out-of-scope path instead of the highest-scoring class.

E. Illustrative Scenario

The following scenario is hypothetical and involves no real learner or institution. A secondary-school learner at a consortium institution types a physics question in Hinglish, using the English name of a law and transliterated Hindi for the rest. The gateway authenticates the learner and resolves the permitted board, subject, grade, and curriculum edition. The intent router returns a concept-explanation intent with high confidence, so the request is not forwarded to BEACON or the credential service. Lexical retrieval matches the named law exactly, dense retrieval surfaces a differently worded passage, and rank fusion places both near the top; the response model answers in the learner's language and cites the passages. Had the evidence been weak, the system would have asked a clarifying question or referred the learner to a teacher.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Later, the same learner applies for a scholarship at another institution that publishes a GPA threshold. The learner's device constructs a proof for that predicate, and the scholarship office learns that the condition holds but not the transcript. Separately, a counselor at the home institution sees an advisory flag from the locally held risk model, which was trained with updates from other sites but never had access to their records. The three interactions share a gateway and an audit log but exchange different data.

VII. SECURITY, PRIVACY, AND GOVERNANCE MODEL

The trust model considers an honest-but-curious federation coordinator, network interception, unauthorized service access, compromised application components, unintended or malformed learner queries, and misuse of academic verification outputs. The architecture does not assume that all institutions, issuers, or model providers are equally trustworthy. It instead makes trust assumptions explicit and limits the information each participant receives. Table VIII lists the principal risks, the architectural mitigations, and the residual limitations.

Risk / concern	Architectural mitigation	Residual limitation
Central reconstruction of learner records	Local data custody; clipped/noisy updates where formally configured; authenticated federation	Model updates may leak information without a valid differential privacy guarantee; secure aggregation [35] may also be needed
Unauthorized learner-data access	Gateway authentication, authorization, purpose limitation, encryption, audit trails	Misconfigured institutional identity or excessive privileges remain risks
Incorrect or hallucinated tutoring response	Curriculum filtering, hybrid retrieval, source metadata, abstention path	Retrieved sources can be incomplete, outdated, or misinterpreted by the generator
Credential disclosure or replay	Zero-knowledge predicate proof, freshness checks, issuer validation, revocation	Trusted setup, issuer compromise, metadata leakage, and circuit bugs remain concerns
Multilingual intent error	Confidence thresholds, out-of-scope routing, human escalation	Low-resource dialects and code-mixed forms may remain underrepresented
Institutional distribution shift	Local evaluation, versioned schemas, institution-wise monitoring	Global aggregation may still disadvantage small or atypical institutions

TABLE VIII
RISKS, ARCHITECTURAL MITIGATIONS, AND RESIDUAL LIMITATIONS

Privacy-enhancing technologies should not be presented as blanket legal compliance. The Family Educational Rights and Privacy Act (FERPA), the Children's Online Privacy Protection Act (COPPA), the General Data Protection Regulation (GDPR), and India's Digital Personal Data Protection Act, 2023 [7]–[10] have different scopes and obligations. Compliance depends on the institutional role, lawful basis, age of learners, notice and consent requirements, retention, processor relationships, cross-border transfers, and operational practices. PREE AI provides design controls that may support data minimization and accountability; it does not certify compliance.

A. Information-Flow Analysis

Table IX summarizes, for each party, what the architecture intends it to learn and what residual exposure remains. The analysis is intentionally conservative: a party is credited with learning anything that is not provably withheld by a stated mechanism.

Party	Intended knowledge	Residual exposure
Coordinator (A1)	Policy-approved updates; schema version	Update-based inference unless secure aggregation and a valid privacy accountant are used
Relying institution (A3)	Proof validity; public inputs	Issuer identity, timing, and repeated-presentation linkage
Network observer (A2)	Encrypted traffic	Message sizes and timing; endpoint metadata
Model provider	Minimum text needed for translation or generation	Query content if an external model is used without safeguards
Peer institutions	Nothing about other sites' learners	Shared model may encode aggregate behavior

TABLE IX
PER-PARTY INFORMATION FLOW AND RESIDUAL EXPOSURE

B. Prompt Injection and Content Safety

Retrieved curriculum text and learner-supplied text are both treated as untrusted data. Instructions embedded in a retrieved passage or an uploaded file must not alter the system prompt, the permitted tools, or the retrieval filters. The generator receives evidence in a delimited block, is instructed to answer only to the extent the evidence supports, and its output is checked for unsupported claims and for age-inappropriate content before it is returned. Because generative models can still produce unsupported statements even when grounded [26], the abstention path and educator escalation remain part of the contract rather than optional features.

VIII. RESPONSIBLE DEPLOYMENT AND GOVERNANCE

Technical controls do not by themselves make an educational AI deployment responsible. Four governance safeguards are treated as part of the architecture rather than as external policy.

A. Educator-in-the-Loop Oversight

Risk scores produced by BEACON are advisory signals for counselors and teachers, not automated decisions. The interface should present the contributing feature groups and the model version, support override and feedback, and avoid ranking individual learners in ways that invite stigmatization. No adverse academic or disciplinary decision should be taken solely from a risk score, and learners or guardians should have a documented route to question a flag.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

B. Fairness and Documentation

Predictive models in education can encode and amplify bias [27]. Evaluation should therefore report error rates by institution, language group, and other legally and ethically permissible groupings, rather than a single pooled metric. Each deployed model and dataset should ship with structured documentation: model cards for intended use, limitations, and subgroup performance [28], and datasheets for provenance, collection method, and consent basis [29].

C. Consent, Purpose Limitation, and Retention

Each data class carries a declared purpose and a retention period, enforced by the gateway policy in Eq. (8). For minors, institutions must determine the applicable consent mechanism; for example, India's Digital Personal Data Protection Act, 2023 [10] contains provisions on processing children's data and parental consent. Stress-related indicators remain disabled by default and require a documented justification and institutional review before activation. Query logs and translation caches are subject to the same retention rules as other personal data.

D. Transparency Toward Learners

Learners should be told when they are interacting with an AI system, which curriculum sources an answer relied on, and when the system abstained because evidence was insufficient. Language preferences and the option to reach a human educator should be available in every supported language.

IX. FEASIBILITY, EVALUATION PLAN, AND DISCUSSION

The proposed architecture is feasible because its principal components have established implementation pathways: federated optimization can be orchestrated across institutional nodes; BM25 and FAISS can operate as local retrieval components; IndicBERT can provide contextual representations for intent classification; and Groth16 libraries can generate and verify proofs for a defined arithmetic circuit. The integration challenge lies primarily in interface design, identity and policy management, schema harmonization, and the validation of cross-layer assumptions.

A. Minimum Feasibility Demonstration

A compact demonstration is sufficient to test whether the architecture's interfaces are coherent. It need not claim a population-level educational effect. A prototype can use two or more isolated institutional environments with separate synthetic or appropriately de-identified datasets, a small curriculum corpus, a labeled set of multilingual/code-mixed intents, and a toy eligibility circuit. The use of synthetic data must be disclosed, and such a demonstration must not be described as a real student cohort. Table X outlines the plan.

B. Implementation Considerations

Several engineering choices determine whether the contracts above can be met in practice; none is claimed as implemented in this paper. *Placement:* embedding, intent, and translation models should run inside the institutional boundary where possible; if a hosted model is used, the gateway should send only the minimum necessary text and the provider's retention terms should be reviewed. *Latency:* lexical and dense retrieval are independent and can be issued in parallel, and the latency of each stage should be measured separately before any end-to-end claim is made. *Key management:* issuer signing keys, gateway credentials, and verification keys belong in hardware-backed or managed key stores with rotation procedures, and proof generation on a learner's device should avoid transmitting the witness to any server. *Observability:* logs should be structured and content-free by default, with content sampling permitted only under documented debugging procedures. *Scaling:* because each institution owns its indexes, capacity grows by adding institutions rather than by enlarging a central index, although the coordinator must tolerate stragglers and dropped participants in federated rounds.

Component	Minimum demonstration	Evidence to report
BEACON	Two isolated clients complete a federated round without exchanging raw records	Message schema, local/global model update, privacy configuration; differential privacy claim only if the accountant supports it
Hybrid retrieval	Compare BM25, FAISS, and fused ranking on a small annotated curriculum query set	Recall@k or mean reciprocal rank (MRR), examples of exact-term and paraphrase retrieval, source traceability
Intent routing	Run a held-out labeled query set through the classifier	Class-wise precision/recall/F1, language composition, out-of-scope behavior
Groth16	Generate and verify a proof for a GPA-threshold circuit	Circuit statement, public/private inputs, proof validity, library and hardware details
Integrated workflow	Demonstrate one tutoring request and one credential request end-to-end	Sequence diagram, service boundaries, logs with sensitive fields removed

TABLE X
MINIMUM FEASIBILITY DEMONSTRATION PLAN

C. Evaluation Questions

A later empirical study can evaluate five questions: (1) Does hybrid retrieval improve curriculum passage ranking over either channel alone? (2) How does intent-routing performance vary across language and code-mixing patterns? (3) What is the utility–privacy trade-off of BEACON under a formally accounted privacy budget? (4) What are proof-generation and verification costs for the selected circuit and hardware? (5) Does the integrated workflow preserve the intended data boundaries under realistic failure and stress conditions? These questions are deliberately component-specific; a single aggregate score would obscure the different objectives of the subsystems.

D. Metrics and Experimental Protocol

A later empirical study should use component-specific metrics and report them with uncertainty. For retrieval, let Q be a set of annotated queries, G_q the set of passages judged relevant to query q, and $\mathrm{Top}_k(q)$ the first k passages returned. Recall at k and mean reciprocal rank are

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

$$\text{Recall@}k = \frac{1}{|Q|} \sum_{q \in Q} \frac{|G_q \cap \text{Top}_k(q)|}{|G_q|} \tag{9}$$

$$\text{MRR} = \frac{1}{|Q|} \sum_{q \in Q} \frac{1}{\text{rank}_q} \tag{10}$$

where rank_q is the position of the first relevant passage. For intent routing over class set $\mathcal{Y}$, per-class precision P_c and recall R_c yield

$$\text{macro-F1} = \frac{1}{|\mathcal{Y}|} \sum_{c \in \mathcal{Y}} \frac{2P_cR_c}{P_c + R_c} \tag{11}$$

reported alongside per-language results, confusion between near classes, and out-of-scope behavior. For BEACON, the cost of privacy is reported as the utility gap

$$\Delta_{\text{util}}(\varepsilon) = M(\theta_{\text{np}}) - M(\theta_\varepsilon) \tag{12}$$

where M is a task metric such as recall at a fixed alert budget, θ_{np} is the non-private model, and θ_ε is trained under an accounted budget ε; the gap should be reported per institution and not only for the pooled model.

Datasets. Candidate resources include an expert-annotated question set built from an approved syllabus for retrieval; an intent benchmark with an out-of-scope class [25] extended with code-mixed examples for routing; and the Open University Learning Analytics Dataset [30], partitioned by module or presentation to emulate institutions, for a federated simulation. Any partition-based simulation should be described as a simulation, not as a multi-institution deployment.

Ablations and statistics. Retrieval should compare BM25 alone, dense search alone, and RRF fusion, with dense-retrieval baselines in the spirit of [31]. BEACON should compare local-only training, centralized training as an upper bound, FedAvg [20], and a heterogeneity-aware variant such as FedProx [32], with and without clipping and noise. Results should be repeated over multiple random seeds and data splits, and reported with confidence intervals; paired bootstrap tests are appropriate for per-query retrieval comparisons. Privacy budgets should be computed with an explicit accountant such as Rényi differential privacy [33] and reported together with the adjacency definition.

Item	Reproducibility requirement
Software	Library and model versions; container image digest
Randomness	Seeds for splits, initialization, sampling, and noise
Hardware	CPU/GPU type; proving and verification environment
Data	Provenance, licence, de-identification method, partitioning rule
Privacy	Adjacency unit, C, σ, sampling rate, rounds, accountant, ε and δ
Credentials	Circuit source, constraint count, trusted-setup transcript, verification key hash
Curriculum	Source list, edition, permission to index, chunking parameters

TABLE XI
REPRODUCIBILITY CHECKLIST FOR FUTURE EVALUATION

This checklist follows general guidance on reproducible machine-learning research [34] and adds items specific to federated training and zero-knowledge circuits.

E. Design Trade-offs

Federated learning reduces the need to centralize raw data but introduces communication overhead, non-IID optimization challenges, and possible leakage from updates. Differential privacy can bound a defined privacy loss but may reduce model utility, particularly for small institutional datasets. Hybrid retrieval improves coverage across lexical and semantic query forms but adds indexing, versioning, and ranking complexity. IndicBERT-based routing may improve handling of code-mixed requests, but its performance depends on representative training data and careful treatment of low-resource language varieties. Groth16 supports compact verification, but trusted setup and circuit maintenance introduce governance costs. PREE AI makes these trade-offs explicit rather than presenting the technologies as cost-free guarantees.

F. Limitations

This paper specifies an architecture and a validation pathway; it does not establish that PREE AI improves retention, GPA, course completion, or counseling outcomes. Such claims require an appropriately designed study with defined outcomes, a suitable comparison strategy, ethics review where applicable, and transparent reporting. The architecture also depends on institutional willingness to participate, compatible feature definitions, reliable identity and credential issuers, and access to curriculum material that may be indexed and redistributed under the applicable permissions. The proposed design does not eliminate all privacy risks, guarantee factual correctness, or establish legal compliance by itself.

Further work should implement a reference prototype, conduct component-level ablations, evaluate language and institution-wise disparities, formalize the privacy accountant and trust assumptions, audit the credential circuit, and conduct a pilot with appropriate institutional governance. A later longitudinal study may investigate educational outcomes, but it should be treated as a separate empirical contribution.

G. Threats to Validity of the Proposed Evaluation

Four threats should be anticipated. *Construct validity:* retrieval relevance and answer correctness depend on annotator judgment, so inter-annotator agreement should be reported. *Internal validity:* simulated institutions created by partitioning one dataset may understate real heterogeneity. *External validity:* results on one board, language set, or grade band may not transfer to others. *Conclusion validity:* small annotated sets and few seeds can produce unstable rankings between methods, which is why interval estimates are required.

H. Phased Roadmap

Table XII organizes future work into phases with explicit exit criteria, so that claims are made only after the corresponding phase is complete.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Phase	Deliverable	Exit criterion
1	Reference implementation of gateway, retrieval, router, BEACON client, and credential circuit	Interface contracts enforced; component tests pass
2	Component-level evaluation and ablations on open or synthetic data	Metrics reported with intervals under the checklist
3	Simulated multi-institution federation with accounted privacy budget	Utility gap and per-site results documented
4	Governance review and limited pilot with consenting institutions	Ethics approval; documented oversight procedure
5	Controlled outcome study	Pre-registered outcomes and comparison design

TABLE XII
PHASED IMPLEMENTATION AND EVALUATION ROADMAP

X. CONCLUSION

PREE AI proposes a zero-trust, institution-sovereign architecture for educational intelligence in which multilingual interaction, curriculum-grounded tutoring, collaborative learner-risk modeling, and selective academic verification are coordinated without requiring a central repository of raw student records. Its contribution is the explicit integration of four distinct trust-sensitive functions: local-first federated analytics, hybrid sparse-dense retrieval, multilingual intent routing, and zero-knowledge credential verification. The paper defines the principal data flows, subsystem responsibilities, mathematical abstractions, workflows, and security assumptions needed to implement this design. Rather than claiming unverified learning gains or universal privacy guarantees, it establishes a technically grounded blueprint and a compact feasibility agenda. Future implementation and independent evaluation can determine the practical utility, privacy–utility trade-offs, and educational impact of the architecture.

Looking ahead, the most immediate step is the reference implementation and the staged evaluation outlined in Table XII, beginning with interface-contract tests and progressing to a small, ethically approved pilot. Because the architecture is modular, institutions can adopt individual subsystems incrementally, for example curriculum-grounded retrieval before federated analytics, without weakening the trust boundaries defined here.

DECLARATIONS

Scope of claims. This paper reports an architecture and an evaluation protocol. It reports no measured performance, no human-subject data, and no institutional records. *Data and code availability.* No dataset was generated. A reference implementation is not claimed in this work.

ACKNOWLEDGEMENT

The authors express their sincere gratitude to the School of Data Science and Forecasting (SDSF), Devi Ahilya Vishwavidyalaya (DAVV), Indore, for providing the academic environment, institutional support, and resources necessary to carry out this work.

The authors also wish to commend the School of Data Science and Forecasting for its vibrant and forward-looking academic culture. The school's programmes in data science, artificial intelligence, and forecasting combine rigorous foundations with applied, problem-driven learning, and its faculty consistently encourage students to take up research questions of real social relevance. The open and collaborative atmosphere of the school, together with its emphasis on the responsible and trustworthy use of data, shaped the privacy-preserving, institution-centered outlook of this work. The authors are proud to be associated with SDSF and thankful for the opportunities it has given them to learn, experiment, and publish.

The authors are particularly grateful to Dr. Govind Maheshwari, Head of Department, SDSF, DAVV, who guided the project and provided academic direction, institutional support, and encouragement. His guidance throughout the research programme is sincerely acknowledged.

The authors also extend their heartfelt appreciation to Dr. Ashok Kumar Gupta, Professor, SDSF, DAVV, who served as mentor for statistics and scientific methodology, and who provided constructive suggestions and continuous support throughout the course of this work.

Md Kaife contributed to this work through web development and the literature review.

Finally, the authors acknowledge everyone who contributed, directly or indirectly, to the development and completion of this work.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

REFERENCES

[1] R. Luckin, W. Holmes, M. Griffiths, and L. B. Forcier, *Intelligence Unleashed: An Argument for AI in Education*. London, U.K.: Pearson, 2016.

[2] NLLB Team *et al.*, "No language left behind: Scaling human-centered machine translation," arXiv:2207.04672, 2022.

[3] J. Gala *et al.*, "IndicTrans2: Towards high-quality and accessible machine translation models for all 22 scheduled Indian languages," *Trans. Mach. Learn. Res.*, 2023.

[4] D. Kakwani *et al.*, "IndicNLPSuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages," in *Findings Assoc. Comput. Linguistics: EMNLP 2020*, 2020, pp. 4948–4961.

[5] K. E. Arnold and M. D. Pistilli, "Course Signals at Purdue: Using learning analytics to increase student success," in *Proc. 2nd Int. Conf. Learn. Analytics Knowl. (LAK)*, 2012, pp. 267–270.

[6] J. P. Campbell, P. B. DeBlois, and D. G. Oblinger, "Academic analytics: A new tool for a new era," *EDUCAUSE Rev.*, vol. 42, no. 4, pp. 40–57, 2007.

[7] Family Educational Rights and Privacy Act of 1974, 20 U.S.C. § 1232g.

[8] Children's Online Privacy Protection Act of 1998, 15 U.S.C. §§ 6501–6506.

[9] Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation), *Off. J. Eur. Union*, vol. L 119, pp. 1–88, May 2016.

[10] Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023), Government of India, Aug. 2023.

[11] A. C. Graesser, P. Chipman, B. C. Haynes, and A. Olney, "AutoTutor: An intelligent tutoring system with mixed-initiative dialogue," *IEEE Trans. Educ.*, vol. 48, no. 4, pp. 612–618, Nov. 2005.

[12] S. Ritter, J. R. Anderson, K. R. Koedinger, and A. Corbett, "Cognitive Tutor: Applied research in mathematics education," *Psychon. Bull. Rev.*, vol. 14, no. 2, pp. 249–255, 2007.

[13] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Model. User-Adapt. Interact.*, vol. 4, no. 4, pp. 253–278, 1994.

[14] C. Romero and S. Ventura, "Educational data mining: A review of the state of the art," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 40, no. 6, pp. 601–618, Nov. 2010.

[15] R. S. J. d. Baker and K. Yacef, "The state of educational data mining in 2009: A review and future visions," *J. Educ. Data Mining*, vol. 1, no. 1, pp. 3–17, 2009.

[16] E. Kasneci *et al.*, "ChatGPT for good? On opportunities and challenges of large language models for education," *Learn. Individ. Differ.*, vol. 103, Art. no. 102274, Apr. 2023.

[17] P. Lewis *et al.*, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 9459–9474.

[18] S. Robertson and H. Zaragoza, "The probabilistic relevance framework: BM25 and beyond," *Found. Trends Inf. Retr.*, vol. 3, no. 4, pp. 333–389, 2009.

[19] G. V. Cormack, C. L. A. Clarke, and S. Büttcher, "Reciprocal rank fusion outperforms Condorcet and individual rank learning methods," in *Proc. 32nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2009, pp. 758–759.

[20] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2017, pp. 1273–1282.

[21] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014.

[22] M. Abadi *et al.*, "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2016, pp. 308–318.

[23] J. Groth, "On the size of pairing-based non-interactive arguments," in *Advances in Cryptology—EUROCRYPT 2016* (Lecture Notes in Computer Science, vol. 9666). Berlin, Germany: Springer, 2016, pp. 305–326.

[24] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with GPUs," *IEEE Trans. Big Data*, vol. 7, no. 3, pp. 535–547, Jul. 2021.

[25] S. Larson *et al.*, "An evaluation dataset for intent classification and out-of-scope prediction," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, 2019, pp. 1311–1316.

[26] Z. Ji *et al.*, "Survey of hallucination in natural language generation," *ACM Comput. Surv.*, vol. 55, no. 12, Art. no. 248, 2023.

[27] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, Art. no. 115, 2021.

[28] M. Mitchell *et al.*, "Model cards for model reporting," in *Proc. Conf. Fairness, Accountability, Transparency (FAT*)*, 2019, pp. 220–229.

[29] T. Gebru *et al.*, "Datasheets for datasets," *Commun. ACM*, vol. 64, no. 12, pp. 86–92, Dec. 2021.

[30] J. Kuzilek, M. Hlosta, and Z. Zdrahal, "Open University Learning Analytics dataset," *Sci. Data*, vol. 4, Art. no. 170171, 2017.

[31] V. Karpukhin *et al.*, "Dense passage retrieval for open-domain question answering," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2020, pp. 6769–6781.

[32] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. Mach. Learn. Syst. (MLSys)*, vol. 2, 2020, pp. 429–450.

[33] I. Mironov, "Rényi differential privacy," in *Proc. IEEE 30th Comput. Secur. Found. Symp. (CSF)*, 2017, pp. 263–275.

[34] J. Pineau *et al.*, "Improving reproducibility in machine learning research," *J. Mach. Learn. Res.*, vol. 22, no. 164, pp. 1–20, 2021.

[35] K. Bonawitz *et al.*, "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2017, pp. 1175–1191.

[36] P. Kairouz *et al.*, "Advances and open problems in federated learning," *Found. Trends Mach. Learn.*, vol. 14, nos. 1–2, pp. 1–210, 2021.

[37] S. Bowe, A. Gabizon, and I. Miers, "Scalable multi-party computation for zk-SNARK parameters in the random beacon model," *IACR Cryptol. ePrint Arch.*, Rep. 2017/1050, 2017.

AUTHOR INFORMATION

Dr. Govind Maheshwari is the Head of Department at SDSF, DAVV, Indore, India. As Head of Department, he guided the project and provided academic direction, institutional support, and encouragement.

Dr. Ashok Kumar Gupta is a Professor at SDSF, DAVV, Indore, India. He served as mentor for statistics and scientific methodology, providing constructive suggestions and continuous support throughout the work.

Aditya Mandloi is an M.Tech (Dual Degree) student in Artificial Intelligence and Data Science at the School of Data Science and Forecasting (SDSF), Devi Ahilya Vishwavidyalaya (DAVV), Indore, India. His contributions to this work include conceptualization, research design, architectural development, and technical analysis. Also Md kaifee co author who significant contributed in paper.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)