Trends and Reviews of Technological Engineering in Education
Vol. 2, No. 1 (2026) | ISSN 3071-2211 | Open Access

RESEARCH ARTICLE

A Permissioned-Blockchain and Verifiable-Credential Protocol for Stackable Engineering Micro-Credentials

Yang Lu[1,*]
[1] Faculty of Business, Information & Human Sciences (FBIHS), Kuala Lumpur University of Science & Technology (KLUST), Malaysia
[*] Correspondence: ziiyuu@gmail.com

Received 13 November 2025 · Accepted 3 March 2026 · Published 30 March 2026

Abstract

Engineering education increasingly uses short, competency-focused learning experiences to address changing technical practices, yet evidence of achievement is often trapped in institution-specific learning systems. A learner may complete safety, instrumentation, programmable-logic-controller, or digital-design work across several providers, while employers and receiving institutions lack a reliable way to verify the issuer, assessment scope, currency, and relationship among those achievements. This article specifies a permissioned-blockchain and verifiable-credential protocol for stackable engineering micro-credentials. Its design separates cryptographic trust and consortium governance from the educational judgment carried by an assessment. Signed credentials are held by learners or issuers; assessment records and identifying information remain off-chain; and a permissioned ledger records issuer authorization, versioned schema and pathway-policy commitments, and auditable governance events. A pathway resolver verifies issuer scope, credential proof and status, then applies explicit rules for outcome coverage, assessment strength, prerequisites, recency, equivalence, and duplicate evidence. The resulting pathway report is institution-local and does not automatically confer course credit, accreditation, or professional licensure. The article gives a data model, issuance and presentation flows, governance controls, a threat analysis, and a synthetic engineering-laboratory scenario. It also defines a conformance and evaluation plan without claiming empirical validation. The protocol treats privacy, appeal, accessibility, and human review as design requirements, and identifies operational questions that a multi-institution consortium must resolve before deployment.

Keywords: Permissioned Blockchain; Verifiable Credentials; Micro-Credentials; Engineering Education; Competency-Based Assessment; Credential Portability

1. Introduction

1.1. The portability problem in engineering learning

Engineering learning spans degree programmes, short courses, employer laboratories, professional development, and digital practice. A learner may demonstrate a narrow but consequential capability—calibrating a sensor, configuring an industrial controller, applying a safety interlock, or validating a building-information model —outside the institution that later needs to interpret it. Micro-credentials can make such learning more visible and modular, but a badge image or a PDF certificate often carries too little machine-readable information to support reliable comparison [1], [2]. A receiving programme must still determine who issued the claim, what the learner did, what evidence supported the award, which outcomes were assessed, and whether the credential remains current. The Council of the European Union's recommendation describes micro-credentials through learning outcomes, workload, assessment, quality assurance, and identity of the learner and issuer [74]; engineering accreditation frameworks likewise focus on demonstrated outcomes and systematic assessment rather than labels alone [3]. Research on teachers' professional-development badges also illustrates recognition and sharing uses; these uses do not by themselves establish engineering pathway equivalence [4].

The technical problem is not simply to place certificates on a blockchain. A ledger can make a record tamper-evident within its governance assumptions, but it cannot establish that a laboratory assessment was valid, that an outcome map is pedagogically defensible, or that two institutions should treat claims as equivalent. Nor does signing a credential make the claim truthful. An interoperable protocol therefore has to join three separately governed functions: (i) a cryptographic proof that a named issuer signed a credential and that the presented data have not been altered [5]; (ii) a trustworthy, reviewable process by which the issuer assessed the learner; and (iii) an explicit policy for how a receiving institution interprets that evidence for a particular pathway. The protocol proposed here keeps those functions distinct while allowing them to exchange versioned, machine-readable information. Existing work on blockchain micro-credentialing and verifiable digital identities provides the technical background for this separation of credential exchange from institutional judgment [6], [7].

1.2. Why stackability needs more than accumulation

A stackable pathway is not a pile of certificates. It is a rule-governed sequence or combination of learning outcomes in which each component contributes evidence to a larger capability profile. Adding credit-hour totals or counting badges can conceal gaps: two credentials may assess the same introductory task, one may be issued for attendance rather than demonstrated performance, or a safety outcome may depend on a prerequisite that is absent. Conversely, a learner may show equivalent competence through a different but acceptable assessment route. Engineering competency frameworks emphasize integrated application, design constraints, experimentation, communication, teamwork, and professional responsibility [8], [9]; a pathway system must preserve those distinctions rather than flatten them into one score. Badge research shows that educational value depends on instructional design and learner context, and reported benefits are not universal [10], [11], [12].

© 2026 TRTEE. Open Access under CC BY 4.0

The protocol uses the term stackable micro-credential for a portable, bounded claim that can be considered by a versioned pathway policy. It does not assume that the credential has a fixed number of academic credits, that all institutions share one taxonomy, or that an automated resolver can make a final award decision. A resolver may report that an evidence set satisfies a configured rule, identify unsatisfied requirements, and provide a trace showing which claim contributed to which outcome. An authorized academic unit remains responsible for decisions about credit, progression, equivalence, or award. This division supports constructive alignment: assessments and learning outcomes remain linked at the issuing institution, while the receiving institution states in advance what evidence it will accept [13], [14].

1.3. Contribution and scope

This article makes four design contributions. First, it defines a layered architecture in which a permissioned ledger supports consortium trust and version control while learner-facing credentials and assessment evidence remain off-chain. Second, it gives a credential profile and pathway-rule model designed for engineering outcomes, performance levels, assessment types, prerequisites, currency, and non-duplication. Third, it specifies issuance, presentation, verification, revocation, and pathway-resolution flows using current verifiable-credential standards as the interoperability base. Fourth, it derives a governance, privacy, and conformance plan for an institutional consortium using permissioned distributed-ledger infrastructure [15], [16].

The work is a technical design specification with an illustrative scenario, not an empirical study or a systematic review. It does not report learner outcomes, deployment measurements, or a validated equivalency table. The architecture is intended to be implementable in stages and to admit more than one credential serialization or ledger client through versioned profiles. Its primary design target is cross-institution engineering education where several recognized issuers want to make bounded learning claims portable without creating a centralized transcript or a public learner ledger. Design-science research provides a methodological basis for distinguishing artifact specification from subsequent demonstration and evaluation [17], [18].

2. Design Basis and Requirements

2.1. Educational claim and evidence model

A credential represents a claim issued by an accountable organization after a defined assessment. For engineering education, the claim should identify the competency or outcome, the context and level at which it was demonstrated, the assessment method, the relevant rubric or standard version [19], [20], and any conditions or expiry rule. These fields do not replace the rubric or evidence record. They provide a compact, structured summary that a learner can present and a verifier can interpret. The issuing institution keeps the underlying work products, assessor notes, identity checks, and appeals record under its own retention and access policies. The credential may carry a concise evidence descriptor or a controlled evidence reference, but it should not publish raw examination materials, student work, or assessor commentary by default.

© 2026 TRTEE. Open Access under CC BY 4.0

An engineering competency is often demonstrated through performance in a constrained context. A learner who can identify a wiring fault in a simulated low-risk environment has not necessarily demonstrated safe intervention on a live installation. A pathway rule must therefore be able to distinguish an outcome label from its assessment conditions [21], [22], [23]. We model a competency claim as a tuple: (outcome identifier and version, level, context, assessment profile, evidence class, issuer scope, issue time, and applicable validity window). The exact vocabulary is defined by a participating academic or professional community; the protocol only makes the vocabulary referenceable and versioned. This is consistent with the emphasis on outcomes and assessment in engineering accreditation criteria and graduate-attribute frameworks [3], [8], [9].

2.2. System requirements

The following requirements define the proposed profile; their satisfaction remains to be tested. Deployments may impose stricter controls. Changes require a versioned profile and governance review, and must preserve the stated privacy and provenance requirements.

ID	Requirement	Protocol behavior
R1	Issuer accountability	Resolve issuer identity, active status, scope, and signing key at the credential issue time.
R2	Portable integrity	Verify a standards-based credential proof and the exact credential schema/profile version.
R3	Learner privacy	Keep personal claims, evidence, and cross-institution learner histories off the shared ledger.
R4	Meaningful stackability	Evaluate outcome coverage, prerequisites, assessment strength, recency, and duplication as explicit rules.
R5	Version-aware interpretation	Bind each pathway result to the mapping and policy versions used in the evaluation.
R6	Contestability	Return a human-readable reason trace and provide a path to challenge a claim or mapping decision.
R7	Operational resilience	Define key rotation, issuer suspension, status publication, recovery, and service-outage behavior.

Table 1. *Design requirements and the system behavior that satisfies them*

2.3. Non-goals and trust boundaries

The protocol is not a global identity system, a replacement for institutional student-information systems, a universal competency ontology, or an accreditation engine. It does not claim that permissioned membership prevents collusion, that ledger consensus determines whether an issuer is educationally competent, or that a cryptographic proof eliminates the need for assessment moderation. A credential’s signature proves integrity under a key-management assumption; it does not prove the issuer used an appropriate rubric. The consortium’s trust registry can state that an issuer is authorized to issue a given credential type, but the academic governance body must still review its assessment design and quality assurance [3], [9].

The ledger trust boundary covers the organizations that operate nodes, approve policy changes, and endorse transactions. A deployment profile must state which institutions participate, what threshold of organizations is required for governance changes, how an organization is admitted or removed, and what fault model the chosen ordering service supports [24], [25]. In Hyperledger Fabric, membership services and endorsement policies provide configurable organization-level identity and transaction authorization; they do not by themselves guarantee Byzantine fault tolerance or independence among endorsers. The protocol therefore treats the ledger as a governed shared registry, not as a self-authenticating source of educational truth. Classical Byzantine fault-tolerant replication

© 2026 TRTEE. Open Access under CC BY 4.0

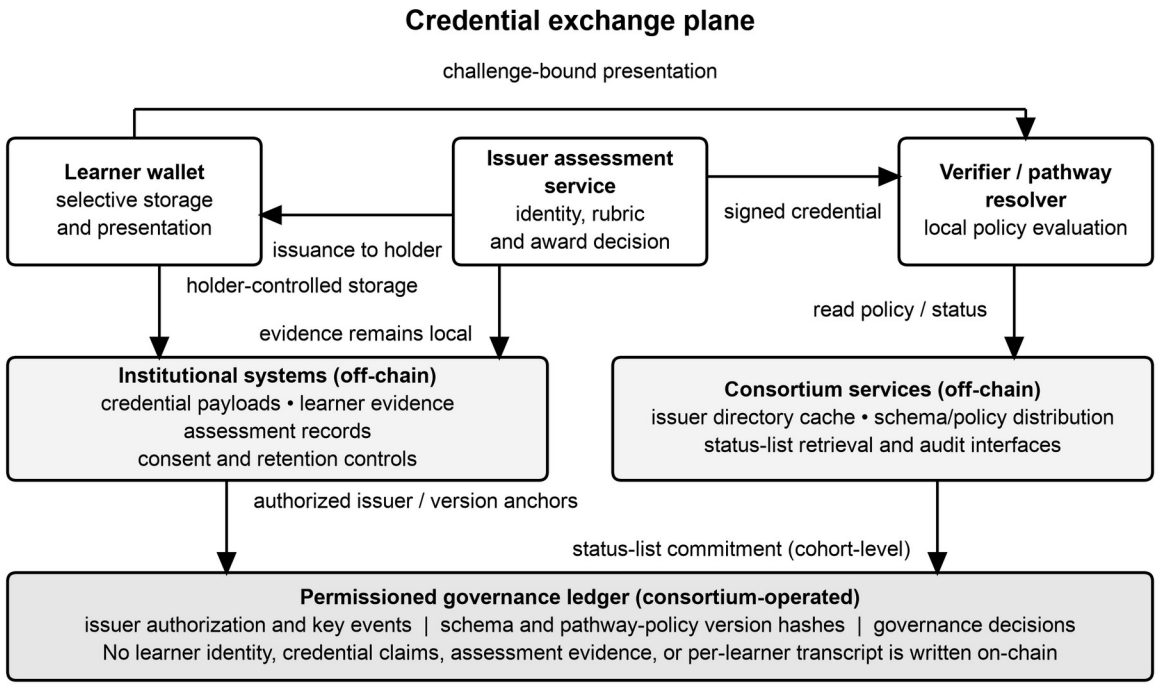
illustrates why fault assumptions and proactive recovery must be specified independently of membership approval [26].

3. Protocol Architecture

3.1. Actors and components

The architecture has five actor classes and four service components. The holder is the learner or an authorized wallet acting for the learner. The issuer is an institution or unit that performs identity checks, assessment, award decisions, signing, and correction. A verifier is a receiving institution, employer, or pathway service that requests selected claims and applies a stated policy. A pathway steward maintains the outcome map and acceptance rules for a named pathway. The consortium governance body approves issuers, schema versions, shared registry policy, and dispute procedures. The roles may be performed by the same institution in a local programme, but separation of duties should be retained for high-impact decisions [27], [28].

The four service components are: (i) a ledger registry for issuer authorization, public-key events, schema and policy commitments, and governance actions; (ii) an issuer service that creates signed credentials and updates status; (iii) a holder wallet that stores credentials and produces minimal presentations; and (iv) a verifier-side resolver that checks proofs, status, trust scope, and pathway rules. Institution-controlled repositories hold evidence and personal records. A status-list service distributes a privacy-conscious revocation or suspension list. Figure 1 shows the separation between the exchange plane, off-chain services, and the governance ledger.



Ledger records govern trust and versioning; credentials and learner data stay under institutional or holder control.

© 2026 TRTEE. Open Access under CC BY 4.0

Figure 1. Credential Exchange and Governance Architecture

3.2. Ledger responsibilities and exclusions

The shared ledger stores information whose common governance benefits from a tamper-evident event history: issuer admission and suspension, signing-key activation and retirement, schema/profile identifiers and content digests, pathway-policy version identifiers and digests, governance votes or approvals, and references to current public status-list resources. The minimum record should contain a stable registry identifier, version, timestamp, authorization context, and a content digest or URI where appropriate. A verifier can resolve the current policy and compare the version digest with the one named by a presentation or pathway request. The underlying idea of committing a document version without publishing its full contents has antecedents in cryptographic time-stamping [29].

The ledger must not store a learner's name, email address, student number, global subject identifier, credential payload, assessment result, evidence pointer tied to one learner, or an individual credential's status transition. Even a salted hash can remain personal or linkable when an attacker can guess the input or correlate the timing and issuer. A hash is not anonymization. If a status list is periodically committed, the commitment should cover a cohort-level list with privacy review and should not expose a per-learner transaction pattern. Credential and evidence storage remain under the holder's and issuer's controlled systems; the ledger carries only the shared trust metadata needed to verify them. Reidentification studies demonstrate why apparently limited metadata can still reveal individuals, while blockchain privacy research highlights correlation risks; these findings motivate a conservative exclusion of learner-level ledger events [30], [31].

Data class	Primary location	Shared ledger?	Reason / control
Credential claims	Holder wallet; issuer copy	No	Contains personal and educational claims; disclose only as needed.
Assessment artifacts	Issuer-controlled repository	No	May contain student work, safety records, or restricted test material.
Issuer identity and scope	Consortium registry	Yes, minimal	Needed to establish who may issue which credential profile.
Schema and pathway versions	Off-chain document plus digest	Digest / version only	Supports reproducible interpretation without publishing learner data.
Key lifecycle events	Consortium registry	Yes	Enables key rotation, retirement, and historical validation.
Per-credential status index	Credential / status service	No individual event	Avoids a durable, linkable learner activity trail.
Cohort status-list commitment	Status resource; optional digest anchor	Possibly, by profile	Use only after linkability and update-timing review.

Table 2. *Recommended data placement and rationale*

3.3. Deployment profile

A first deployment can use a permissioned Fabric network with one organization per institution. Membership Service Providers authenticate organizational identities. Validating peers check endorsements against the applicable endorsement policy [82]; contract logic separately enforces caller authorization and governance rules. Private-data collections can limit visibility of consortium metadata, but introduce replication, retention, backup, and governance obligations. They do not justify placing learner records on the ledger. The recommended design uses off-chain signed credentials and a consortium trust registry [32], [33], rather than private on-chain

© 2026 TRTEE. Open Access under CC BY 4.0

transcripts. A different permissioned ledger may be substituted if it can provide equivalent organization-scoped authorization, auditable governance events, durable version commitments, and a documented fault model [24], [25]. The registry's contract logic also requires explicit authorization and independent review rather than relying on consensus alone [34].

4. Credential and Registry Data Model

4.1. Credential profile

The credential profile uses the W3C Verifiable Credentials Data Model 2.0 [75]. DID Core provides identifier and verification-method infrastructure, not an educational identity policy [76]; a learner need not publish a global identifier. Open Badges 3.0 may supply achievement semantics [77], but each deployment must pin compatible context, schema, and proof versions and test interoperability. The label "3.0" alone does not establish compatibility between all revisions or credential data models. Engineering-specific fields require a versioned extension. Surveys identify wider identity-management, trust, interoperability, and adoption challenges [35], [36].

This proposed profile requires context, type, issuer, subject claims, a versioned schema, an approved embedded proof or secured envelope, and a status reference. It also requires a signed issuance-time extension, issuedAt, defined in that profile; this is not a VC 2.0 core property. Optional validFrom and validUntil delimit validity, not independently proven signing time [75]. Claims identify outcome versions, performance level, context, and assessment profile, with criteria available under appropriate access controls. Selective disclosure must use a declared proof mechanism binding the disclosed claims to the signed credential. Removing JSON fields is insufficient. Anonymous-credential constructions require separately supported cryptographic profiles [37], [38].

Object	Required fields	Purpose
Credential	@context; type; issuer; issuedAt (profile); subject; schema; status; proof/envelope	Portable signed assertion with declared data model and validation method.
Achievement claim	outcomeRef/version; level; context; assessmentProfile; evidenceClass	Describes what was assessed and under which bounded conditions.
Issuer registry entry	orgId; active status; credential scopes; key set; effective interval	Constrains who may issue a profile and which key was valid at issuance.
Outcome map	source outcome/version; target outcome/version; relation; evidence constraints; approver	Makes equivalence explicit, reviewable, and time-versioned.
Pathway policy	required nodes; prerequisites; alternatives; recency; non-duplication; policy digest	Defines how a set of claims is evaluated for one named purpose.
Status resource	list id; purpose; version; publication time; integrity proof	Allows suspension or revocation checks without ledgering each learner event.

Table 3. *Core credential and registry fields in the proposed profile*

© 2026 TRTEE. Open Access under CC BY 4.0

4.2. Competency vocabulary and evidence profile

Outcome identifiers should be stable within a defined vocabulary and accompanied by a version. An outcome record can point to a local curriculum outcome, a programme-level graduate attribute, or a professional framework element, but it should never imply equivalence merely because two strings share a label. The mapping object must declare its relation type, such as exact match, narrower-than, broader-than, partial overlap, or accepted alternative. The mapping owner, review date, supporting rationale, and approval event are required metadata. A pathway resolver may accept only selected relation types for a specific purpose.

Assessment profiles identify methods such as observation, practical work, diagnostics, design explanation, or safety demonstration; rubric use [39], [40], [41]; individual or team work; and identity checks. Claims state profile identifier/version, quality-assurance status, and evidence class; issuers retain the full specification and restricted artifacts. Where recency or learning order matters, the proposed profile also requires issuer-attested assessment-event timing, its precision and time zone, and the applicable comparison rule. Reissuing the same assessment claim cannot reset its age. Issuance and credential validity dates are distinct. Missing timing, or uncertainty that prevents a required comparison, yields indeterminate.

4.3. Versioning and backward interpretation

Schemas, assessment profiles, outcome maps, pathway rules, issuer scopes, and proof profiles are versioned. Interpret signed claims under their declared schema and assessment versions; apply the receiving policy and mappings selected for the evaluation. Historical issuer authorization does not replace current compromise and status checks. Retain authorized input claims, evaluation time, verifier configuration, policy and map versions, registry state, and verified status snapshots to reproduce a past result. Digests alone are insufficient. Reproduction is limited by permitted retention; a new decision requires a fresh evaluation under its applicable policy.

Versioning also limits silent semantic drift. A programme may change what ‘PLC diagnostics level 2’ means as equipment, safety practice, or curriculum changes. Reusing an identifier without a version would make an older credential appear to attest the newer outcome. The protocol therefore requires a new version whenever an outcome’s scope, evidence threshold, or material context changes. A governance review can define equivalence between versions, but the relation is recorded as a separate mapping decision rather than embedded in a mutable label.

5. Issuance and Credential Maintenance

5.1. Issuance protocol

Issuance begins with a learner’s enrollment in an assessed activity and ends with a signed credential held by the learner. The issuer establishes the applicable schema and assessment profile before assessment begins. The learner is informed what claim will be issued, what identity checks are required, what evidence is retained, how long it is retained, and how an appeal or correction can be requested. The assessment service records the rubric and

© 2026 TRTEE. Open Access under CC BY 4.0

evidence under institutional access controls; an authorized academic decision-maker confirms the award. The signing service then creates the credential using a key authorized for the issuer and credential scope. Issuance is logged in the issuer's system, while the ledger records only registry or governance events needed by other consortium members. The award decision should remain distinguishable from formative feedback and support for self-regulated learning [42], [43].

A concrete issuance sequence is: (1) resolve the active credential schema, assessment profile, and issuer scope; (2) verify the learner's identity using the institution's approved process; (3) assess the work against the published criteria; (4) record the decision and any review or appeal; (5) construct the minimum credential claims; (6) sign with the active issuer key and attach an approved status reference; (7) deliver the credential through an authenticated wallet channel or download; and (8) provide a human-readable receipt and correction contact. The issuer should not write a hash of each learner's credential to the ledger as an issuance receipt. Signed dates assert issuance time; independent timing needs accepted timestamps or corroboration, without learner-level ledger entries [29].

5.2. Key lifecycle and issuer authorization

Issuer keys are distinct from node-administration keys. Registry entries identify issuer, scope, key, proof suite, activation and retirement intervals, and compromise status. Institutional key management requires role separation [27], [28], access logs, recovery, and emergency rotation. Retirement need not invalidate earlier signatures, but current compromise policy still applies. An attacker with a compromised key can backdate a credential: a signed issuedAt or validFrom cannot establish pre-compromise issuance. Governance must specify accepted independent timing evidence, affected intervals, reissuance, and review. Unresolved historical timing yields an indeterminate result, not automatic acceptance.

Issuer authorization should be more specific than a binary 'member' flag. A member may be permitted to issue a safety-training credential but not a degree-level outcome or an industry-standard certification. The registry therefore records scopes as a set of credential profiles and outcome families, along with any quality-assurance conditions and renewal dates. The consortium may suspend a scope, suspend the whole issuer, or retire an organization. The verifier must evaluate the authorization state at the relevant issuance time and, when policy requires it, the current state at presentation. The two checks answer different questions: whether the issuer was authorized when it made the claim and whether the issuer is currently trusted for the requested use [35], [44].

5.3. Correction, replacement, suspension, and revocation

Educational records sometimes require correction because of clerical error, successful appeal, issuer-key compromise, or an assessment decision later found invalid. A credential is not silently edited. The issuer issues a replacement with a new identifier and versioned correction reason, updates the old credential's status, and preserves a restricted audit trail under its records policy. A verifier should be able to distinguish revocation for fraud or invalid assessment from ordinary expiry, replacement, or administrative suspension, but the public status channel should disclose only the minimum status category required by the

© 2026 TRTEE. Open Access under CC BY 4.0

profile. Detailed reasons remain available through an authorized appeal or verification process.
W3C Bitstring Status List uses authenticated list credentials [78]. Check the authorized status issuer, credential-to-list binding, proof, purpose, index bounds, validity, and policy freshness. Revocation is irreversible; suspension can be lifted through an authorized status update. Neither revoked nor currently suspended credentials support completion. The ttl refresh hint does not extend list validity. A permitted cached list bounds status delay; it cannot prove that no newer list exists. Missing, stale, or unavailable required status yields indeterminate. Record manual decisions separately. Revocation does not itself rescind an academic award. Accumulators require a separate proof profile [45].

6. Presentation, Verification, and Pathway Resolution

6.1. Challenge-bound selective presentation

A verifier requests a purpose and minimum claims, identifying the audience, expiry, trust scope, status policy, and pathway version. The wallet displays the request and recipient. The verifier binds a fresh unpredictable nonce to the session and atomically consumes it on acceptance; expiry alone does not stop reuse within the window. The presentation must bind that challenge and audience to a holder key or other binding authorized by the issuer credential [81]. A proof from an arbitrary wallet key does not establish ownership of a copied credential. Holder control and learner identity assurance are separate checks. Request integrity and supported disclosure profiles must also be verified [35], [37].
The W3C credential model is separate from its securing mechanism. A deployment declares supported Data Integrity [79] or JOSE/COSE profiles [80], including any selective-disclosure and holder-binding requirements. In an SD-JWT-based exchange, disclosure verification alone is insufficient when key binding is required [81]. The verifier must establish that disclosed claims belong to the issuer-signed credential [37], [38]. Derived proofs and zero-knowledge extensions need separate cryptographic review, interoperability testing, and maturity assessment. This design proposes profile requirements; it does not establish conformance or security of an implementation.

6.2. Verification pipeline

The verifier returns a result for each stage in Table 4. First check request integrity, audience, freshness, one-use nonce, and credential-authorized holder binding. Second verify credential proofs and schema. Third resolve issuer keys and scopes, historical authorization, and applicable current compromise policy. Fourth check validity and authenticated, fresh status. Fifth interpret outcome and map versions, assessment profiles, and evidence classes. Sixth pass only eligible claims to pathway resolution. Unsupported formats or unavailable mandatory checks yield indeterminate; failed security or trust checks yield invalid. A valid claim can still fail educational eligibility: a cryptographic failure

© 2026 TRTEE. Open Access under CC BY 4.0

means the data cannot be trusted as a signed claim, whereas a policy mismatch means the claim does not satisfy the requested pathway [7], [35].

Stage	Check	Failure response
1. Request	One-use nonce, audience, expiry, holder binding	Reject stale, mismatched, or malformed presentation.
2. Proof	Credential proof and schema conformance	Invalid proof: invalid; unsupported format: indeterminate.
3. Trust	Historical scope/key; current compromise policy	Mark untrusted; provide issuer-scope reason.
4. Status	Validity; bound, authenticated and fresh status	Revoked/suspended: exclude; unknown: indeterminate.
5. Semantics	Outcome/map versions; assessment and evidence	Return unsupported or unmapped claim; do not guess.
6. Pathway	Coverage, prerequisites, recency, duplicate use	Return unsatisfied rules and an explanation trace.

Table 4. *Verifier checks and expected handling*

6.3. Formal pathway semantics

A versioned pathway graph G = (V, E) has competency nodes and typed prerequisite edges. Reject cyclic prerequisites; express alternatives separately as Boolean or threshold rules. Policy P specifies issuers, assessment profiles, evidence classes, mappings, levels, recency, context, exclusions, and purpose. Required learning order needs assessment-time evidence, not credential signing order. These proposed requirements express a receiving institution’s evidence-acceptance policy; they do not establish curriculum equivalence.

Let C be presented credentials, Verify(c,t) their verification results at time t, and Map(c,P) the permitted mapped claims. A node requires eligible proof, issuer scope, status, assessment, performance, and mapping. Completion requires a joint allocation of evidence to all mandatory nodes, respecting prerequisites, alternatives, recency, context, exclusions, and explicit reuse limits. Independent node matches are insufficient if they reuse evidence impermissibly. Credentials must also be bound to the same learner through the approved identity-assurance process; control of unrelated holder keys alone does not prove this. Avoid a global tracking identifier for that check.

Duplicate control may require an issuer-scoped assessment-event identifier or approved attestation; an evidence-class code alone is not unique. Replacements retain the event relationship. Policy defines permissible reuse across outcomes without creating a cross-institution tracking identifier. An unavailable mandatory uniqueness check yields indeterminate and authorized review, never completion by caveat. Changing reuse requirements requires a newly approved policy version and a fresh evaluation; it cannot waive a check in the current evaluation. Record node support, unmet rules, policy and mapping versions, and unresolved checks.

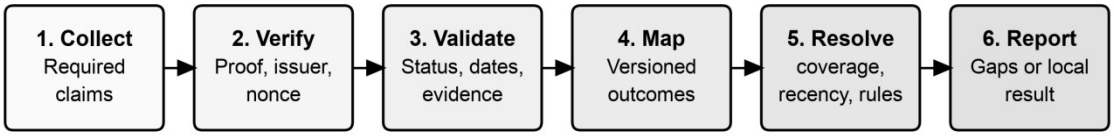
6.4. Pathway result and auditability

Return invalid for failed required security or trust checks; else indeterminate for unresolved mandatory checks; else incomplete for known gaps; else complete. Record credential exclusions and reasons. The signed result names purpose, policy digest, evaluation time, status versions, and minimal references. Re-evaluate before a later decision. Do not copy the full credential

© 2026 TRTEE. Open Access under CC BY 4.0

set into a new central learner record unless the receiving institution has a separate, explicit basis and retention policy. A learner-facing explanation translates rule identifiers into concrete actions—for example, ‘supervised safety demonstration is required’—rather than exposing only a machine code.

The audit trail is purpose-limited. The receiving institution may retain an evaluation record needed to justify its own decision, but should avoid placing that record on the shared ledger. A high-stakes decision should preserve a reviewable trace of the rules used and a path for the learner to challenge an incorrect mapping, stale status, or identity mismatch. Reproduction requires the permitted inputs, configuration, time, and snapshots specified in Section 4.3; policy and registry versions alone are insufficient. Unrelated credentials need not be retained. The result is evidence for institutional judgment, not an automated degree award or professional license. Assessment research on evaluative judgment reinforces the importance of intelligible criteria and human judgments of quality [46].



A credential can be cryptographically valid while failing a pathway rule (for example, wrong scope, stale evidence, insufficient assessment, or duplicate coverage).

The resolver returns a result to the requesting institution; it does not publish a learner profile to the ledger or confer academic credit automatically.

Figure 2. Pathway Resolution and Educational Rule Validation

7. Governance and Consortium Trust

7.1. Membership and role separation

A consortium charter defines admission criteria, organizational responsibilities, governance quorum, conflict-of-interest handling, audit rights, service levels, incident reporting, and dispute resolution [16]. Admission is a substantive educational decision: an institution must demonstrate authority to issue the credential types it requests, identify accountable academic officers, document assessment and quality-assurance processes, and agree to privacy, security, retention, and accessibility requirements. Network membership alone is insufficient. The registry records the scope and effective dates of approval, while the charter contains the criteria and review process.

Role separation reduces the risk that one compromised account can approve an issuer, change a pathway rule, and issue credentials under the new rule. At minimum, the governance process separates proposer, reviewer, approver, and network operator roles [27], [28]. For a high-impact pathway, the change proposal should include a rationale, impact analysis for learners with older credentials, mapping changes, transition rules, and an appeal route. The ledger can make the approvals and version transitions tamper-evident; the actual deliberation record may remain in a governed document repository with an on-ledger digest.

© 2026 TRTEE. Open Access under CC BY 4.0

7.2. Schema and pathway-policy change control

A pathway steward proposes a new or revised policy using a structured change package. The package specifies target outcomes, prerequisite logic, accepted issuer scopes, assessment profiles, equivalency relations, recency windows, non-duplication requirements, and the rationale for each rule. Reviewers include engineering educators, assessment specialists, student representatives where appropriate, and receiving-program decision makers. Before activation, the change is tested against synthetic cases including valid, borderline, incomplete, duplicate, expired, and adversarially malformed presentations [47]. A new policy version receives a digest and activation time; the former version remains retrievable for historical evaluations.

Governance should not retroactively lower an evidence threshold for learners who earned credentials under a different standard without stating that decision. A policy can preserve old credentials, require a bridging assessment, recognize an approved equivalence, or narrow acceptance prospectively. Each choice has educational consequences and should be visible in the pathway's published interpretation. For a change that affects safety-relevant outcomes, the steward may require review by an external professional or regulatory body, but the protocol does not itself grant that body authority.

7.3. Appeals, correction, and accountability

A learner should be able to ask which claim failed, which versioned rule was applied, whether the issue was proof, status, issuer scope, or educational sufficiency, and how to request correction. The issuer handles challenges to assessment and credential accuracy; the pathway steward handles challenges to outcome mapping or acceptance policy; the verifier handles implementation errors in an individual evaluation. A shared protocol makes those boundaries explicit so that a learner is not sent between institutions without an accountable owner. Appeals should not require broad disclosure of the learner's other credentials. Where an incomplete pathway leads to further learning, explanation design can draw on research on feedback processes and student feedback literacy [48], [49].

The consortium should publish a registry of active policy versions, issuer scopes, contact channels, and governance decisions that affect credential interpretation. Where the public record could reveal confidential institutional operations or create security risk, detailed records can be access-controlled while a digest and high-level decision remain discoverable. Accountability is not equivalent to public exposure: privacy-preserving governance provides enough information to understand who made a rule and when, while withholding learner-level data and restricted assessment material.

8. Security, Privacy, and Threat Model

8.1. Threat assumptions

The threat model assumes that attackers may steal a learner's credential file, replay a presentation, phish a verifier, compromise an issuer account, exploit a signing key, tamper with an off-chain repository, submit a forged outcome mapping, collude with one consortium member, or infer learner activity from registry and status requests [50]. A participating institution may also misconfigure a pathway resolver, retain excessive data, or

© 2026 TRTEE. Open Access under CC BY 4.0

publish a weak assessment profile. The protocol assumes that cryptographic primitives are implemented correctly, that the ledger has enough independent operators for its stated fault model, and that institutions enforce their own access controls. It does not assume that all consortium members are honest or that network consensus can correct a false academic judgment. The model also covers vulnerabilities in registry contract logic, because consensus does not eliminate software defects [51].

Threat	Primary control	Residual risk / response
Credential alteration or forgery	Approved proof suite; issuer-scope and key validation	Compromised issuer key can produce apparently valid claims; rotate, suspend, and investigate.
Replay or presentation substitution	One-use nonce; audience; expiry; credential-bound holder proof	Verifier or wallet bugs remain possible; test and monitor protocol implementation.
Learner correlation	No on-chain subject data; pairwise identifiers; minimal disclosure	Timing, issuer, or status queries may still correlate activity; use caching/proxy and limit logs.
False or weak educational claim	Assessment-profile registry, moderation and issuer review	A signature cannot validate pedagogy; governance and appeals remain necessary.
Policy manipulation	Versioned change request, role separation, multi-organization approval	Coalition or governance capture remains possible; publish audit and dispute procedures.
Status-service outage or stale list	Signed lists, freshness checks, monitored distribution, defined fallback	Availability may delay decisions; report indeterminate, do not silently accept.
Ledger collusion or operator failure	Independent organizations, endorsement thresholds, backups, incident plan	Permissioned consensus follows its configured fault model; cross-organization collusion is not eliminated.

Table 5. *Threats, controls, and residual risks*

8.2. Privacy by data minimization

Privacy begins with minimum claims for a defined purpose, a clear wallet request, and holder choice. Do not publish stable learner identifiers to the ledger. Pairwise identifiers may reduce some linkability [32], [52], but a reused credential identifier, signature, key, or status index can still correlate presentations. They do not guarantee unlinkability. Document loss, recovery, duplicate handling, and accessibility. Request a narrow claim or supported derived statement when sufficient, and evaluate correlation across the complete presentation, status, and logging flows.

Contextual integrity asks whether the sender, recipient, subject, attribute, and transmission principle fit the educational setting [53]. Status queries may expose presentation activity. Verifier-side caching, shared cohort lists, proxies, or holder retrieval can reduce direct issuer observation while retaining freshness checks. An index disclosed with a credential can correlate repeated use; a sparsely populated or uniquely assigned list can also weaken privacy. Padding alone does not establish anonymity [78]. Proxies introduce trust and logging risks. Document the chosen trade-offs and threat assumptions [31], [35].

8.3. Cryptographic agility and zero-knowledge extensions

Cryptographic algorithms and proof formats age. The consortium must approve a limited proof suite, specify key sizes and curves, maintain a migration timeline, and test verifier behavior for unknown algorithms. The credential identifies its suite, and the registry

© 2026 TRTEE. Open Access under CC BY 4.0

constrains which suites are accepted for each issuance period. A verifier must fail closed on an unknown or disallowed proof suite while preserving the signed artifact for later review. Migration must account for credentials that remain valid after an algorithm is deprecated, including whether an issuer can reissue them without repeating the original assessment. Zero-knowledge proofs, first formalized within interactive proof systems, may eventually allow a learner to prove a predicate—such as ‘has completed at least three distinct required outcomes’—without disclosing each underlying credential [54], [55], [56]. Such a design could reduce unnecessary exposure, but it introduces assumptions about circuit correctness, proof-system security, revocation handling, and interoperability [38], [45]. The protocol treats zero-knowledge presentation as an optional extension, not a baseline requirement. It must undergo an independent cryptographic review, define how issuer trust and status are proved, and provide human-readable explanations for the predicate. A privacy feature that is too complex for learners or verifiers to understand can shift risk rather than remove it.

8.4. Ledger security and availability

Permissioned network security requires administrative controls as well as consensus. Separate responsibilities for node credentials, contract lifecycle, endorsement policy, orderers, backups, monitoring, and incident response. One compromised peer must not authorize issuer scopes. Test organization removal, key rotation, node loss, snapshot restoration, and policy rollback through a newly authorized version, preserving history [26], [50]. The proposed profile defines authenticated version ordering for each registry or status resource and rejects rollback below the latest trusted version retained for it. Outages do not relax validity or freshness limits. Recovery checks authenticated checkpoints; unresolved version order or freshness blocks dependent acceptance.

Availability, confidentiality, and integrity are different service objectives. The ledger may be highly available while the issuer’s assessment archive is unavailable; a valid proof may exist while an issuer scope is suspended; and a pathway resolver may be unavailable even though all credentials remain valid. User interfaces should report these conditions separately and avoid a single green ‘verified’ badge that collapses them. Operational telemetry should use aggregate service metrics and avoid logging full presentations or subject identifiers unless a justified incident process requires them.

9. Engineering-Education Integration and Illustrative Scenario

9.1. Curriculum and assessment design

The protocol begins with curriculum design, not ledger deployment. A programme identifies an observable engineering outcome, chooses an authentic assessment, states the performance threshold, and determines what contextual information is necessary for another institution to interpret the claim. An outcome such as ‘diagnose a control-system fault’ should specify whether the learner reasons from a schematic, interacts with a simulator, measures a physical signal, or performs work under supervision. The assessment must align with the learning outcome and the intended use. A credential is only as meaningful as the assessment and criteria it summarizes; cryptographic portability cannot compensate for a poorly aligned assessment [13], [19]. Evidence from engineering and STEM education supports considering active and inductive learning when designing the

© 2026 TRTEE. Open Access under CC BY 4.0

assessed activity [57], [58], [59]. Assessment feedback should help learners understand performance criteria and plan subsequent learning [60], [61]. These findings inform educational design; they do not validate the credential protocol.

Outcome maps should be co-designed by the issuing and receiving communities. The issuer's goal is to describe what its assessment establishes; the receiving institution's goal is to state what evidence supports its pathway. Mapping meetings should include technical faculty, laboratory staff, assessment experts, and where relevant industry partners. An employer may explain workplace context, but should not unilaterally define academic equivalence. A map can accept a narrow credential as partial evidence for one outcome while still requiring a local practical examination. This preserves curricular autonomy and avoids interpreting common vocabulary as automatic equivalence. Critical analyses of micro-credentials warn that granular awards can fragment curricula and privilege immediate workplace demands over broader educational purposes [62], [63].

9.2. Worked scenario: industrial controls laboratory readiness

Consider a synthetic pathway titled Industrial Controls Laboratory Readiness shared by two engineering schools and a continuing-education provider. The pathway requires three outcome nodes: (A) identify hazards and apply a defined lockout procedure in a supervised laboratory; (B) configure and verify a programmable logic controller (PLC) interlock using a stated design specification; and (C) diagnose a sensor or actuator fault using measured signals and explain the resulting evidence. The pathway policy requires direct practical assessment for A and B, permits either practical bench work or a validated hardware-in-the-loop exercise for C, and requires that the safety claim be current under the receiving programme's published recency rule. These are scenario rules, not universal engineering requirements.

Provider 1 signs a PLC interlock credential after individual rubric-based bench assessment and assessor sign-off. Provider 2 issues a safety credential after observed lockout practice. A learner presents only the achievement claims and proof of possession needed for the pathway. The resolver confirms issuer scope and status, maps each outcome version to the pathway version, checks assessment profiles and dates, and identifies node C as missing. The explanation identifies missing fault-diagnosis evidence and accepted assessment options. After completing a hardware-in-the-loop fault diagnosis at a third issuer, the learner presents all required claims in a fresh challenge-bound exchange. Only after rechecking current status, recency, prerequisites, and evidence reuse can the resolver report all nodes satisfied under the named policy version. It does not award university credit or certify the learner as competent for unsupervised industrial work.

The scenario illustrates why the protocol treats evidence quality, recency, context, and pathway purpose as separate fields. A badge count would show three achievements but not whether the required safety practice was observed. A sum of hours would not establish whether PLC configuration was assessed individually. An unversioned mapping could silently change after the learner completes the work. A shared ledger that stored the learner's pathway result would expose a durable education history without being necessary for verification. The proposed design instead allows each institution to make a bounded claim and each verifier to evaluate it for a declared purpose.

© 2026 TRTEE. Open Access under CC BY 4.0

9.3. Accessibility and learner agency

Portability should not require learners to understand cryptographic key management. Wallet interfaces need plain-language consent screens, accessible presentation of issuer and claim information, recovery options [44], and a clear route for requesting a credential in a non-digital format where needed. A learner who changes phone, loses a key, or cannot use a wallet should not lose the underlying academic achievement. Issuers retain a route to reissue credentials after identity verification, and institutions retain an appeal and records process. Recovery must avoid turning the issuer into a silent observer of every presentation or creating a single central wallet database.

The protocol also avoids making possession of a particular device or digital identifier a learning outcome [64]. Where assessment requires identity assurance, that assurance is recorded as an assessment condition and reviewed for proportionality. Learners should be told which data are disclosed to a verifier, whether a presentation can be retained, and how long the receiving institution plans to keep it. Consent interfaces cannot resolve every power imbalance—students may feel compelled to disclose [53], [65]—but data minimization and purpose limitation reduce unnecessary collection. A future deployment should include accessibility testing and learner consultation as acceptance criteria, not as post-launch enhancements.

10. Deployment Profile and Conformance Evaluation

10.1. Implementation sequence

A proposed four-phase sequence starts with governance, registry, vocabulary, and one synthetic pathway. Phase 2 tests proof and status checks, correction, suspension, key rotation, compromise, recovery, and incident procedures using synthetic records; blocking conformance defects must be resolved. Phase 3 permits a small real pilot only after technical, privacy, accessibility, and governance gates pass. Phase 4 expands participation after pilot evidence and workload review. Each phase has a stop and rollback route. A permissioned ledger requires a shared governance and audit need; a signed registry may suffice for a small consortium [66], [67], [68].

The profile pins data model, schema, vocabulary, proof and holder binding, request formats, nonce lifecycle, status and version-order rules, registry API, canonicalization, and failure codes. Unsupported mandatory features yield indeterminate. Synthetic fixtures test tampering, copied credentials with unrelated keys, replay, expiry, suspension, revocation, stale or rolled-back lists, unknown issuers, compromise and backdating, mapping gaps, partial disclosure, cyclic prerequisites, and prohibited reuse. Include uncertain assessment timing and reissuance that must not renew recency. Matching results require identical inputs, policy, time, configuration, and snapshots [47], [69].

10.2. Conformance and evaluation plan

Because this article specifies a design and does not report a deployed system, the following evaluation plan is a proposal rather than a result. Technical evaluation should measure

© 2026 TRTEE. Open Access under CC BY 4.0

proof-verification correctness, registry consistency, status freshness, resilience to key rotation and service interruption, data minimization, and cross-vendor interoperability. Educational evaluation should assess whether pathway rules accurately reflect programme expectations, whether evidence descriptors are intelligible to faculty and learners, and whether the resolver's explanation supports a fair human decision. Operational evaluation should record administrative burden, correction and appeal rates, and the frequency of indeterminate outcomes. These measures should be defined before pilot data are collected [70], [71].

Test family	Representative case	Expected acceptance evidence
Proof integrity	Change one signed claim or proof byte	Verifier rejects or reports invalid proof; no pathway completion.
Issuer scope	Valid signature from issuer outside requested scope	Proof may pass; trust-scope check fails with a specific reason.
Status	Revoked/suspended; stale or missing list	Revoked/suspended: no completion; unknown: indeterminate.
Challenge security	Replay presentation with prior nonce or other audience	Reject used nonce, wrong audience, or missing holder binding.
Rule semantics	Missing prerequisite, duplicate evidence, expired practical claim	Resolver returns incomplete with the violated rule and trace.
Privacy	Inspect ledger, logs, and verifier storage after presentation	No learner claims or full presentation written to shared ledger; logs minimized.
Governance	Rotate key, suspend issuer, activate policy v2	Historical and current states resolve consistently and are auditable.
Interoperability	Exchange between independently implemented wallets and verifiers	Same profile, inputs, policy, time, configuration and snapshots produce the same result.

Table 6. *Suggested conformance suite and acceptance evidence*

10.3. Metrics and decision gates

The consortium should set service targets after baseline testing rather than copy arbitrary thresholds from another deployment. Technical metrics include proof-verification latency distribution, status-list age, registry propagation delay, outage duration, false acceptance and false rejection rates, and recovery time after key rotation. Educational metrics include assessor agreement for the underlying rubric, pathway-rule agreement between faculty reviewers, rate of human overrides, appeal outcomes, and learner comprehension of disclosure requests. Equity review should examine whether learners from different institutions, languages, disability contexts, and device environments experience different rates of credential presentation failure. Learning gains and learner-reported perceptions of learning should be measured separately, because perceived learning need not track tested learning [72]. Evidence that teaching design can affect STEM achievement gaps motivates evaluating the assessed learning activity as well as unequal access to credential presentation [73].

A pilot should not advance simply because the cryptographic checks pass. The decision gate asks whether the pathway's rules are defensible, whether a learner can correct and appeal an error, whether a verifier can explain an indeterminate result, whether the consortium can respond to compromised keys, and whether participating institutions can sustain the workload. If the pilot finds that the outcome map is disputed or the assessment profile is too weak to support the intended decision, the correct response is to revise the educational rule or narrow the use case—not to add more ledger transactions [67], [71].

© 2026 TRTEE. Open Access under CC BY 4.0

11. Discussion and Limitations

11.1. What the protocol enables

The proposal offers a way to make bounded engineering achievements verifiable and composable without centralizing a learner's cross-institution history. It makes the semantics of stacking explicit: pathways are versioned policies over outcomes and evidence, not numeric sums. It also assigns responsibility to the organizations best placed to perform each task. Issuers own assessment and correction; consortium bodies govern issuer scopes and shared schemas; pathway stewards state acceptance rules; verifiers explain the result; and learners control presentation. A ledger contributes a shared, auditable registry when institutions need common governance, but signed credentials can remain useful if the ledger is unavailable or replaced.

The architecture can support a distributed ecosystem in which local curricula remain different while selected evidence can travel. This is valuable in engineering because capabilities are context-sensitive and assessment conditions matter. A common protocol does not require one global taxonomy or a single accreditor. Instead, it allows participants to make explicit where their claims align, where they are partial, and where no mapping has been approved. This stance avoids the false choice between total standardization and non-portability.

11.2. Limitations

Several limitations remain. First, outcome mapping is a normative academic judgment and cannot be solved by cryptography. Second, assessment quality is only as strong as the issuer's design, moderation, and accountability. Third, selective disclosure is not uniform across credential proof formats, and advanced privacy proofs can impose implementation and usability costs. Fourth, status checking may itself leak activity, especially when verifiers query issuer-hosted endpoints. Fifth, a permissioned ledger introduces governance and maintenance burdens that may exceed its value for a small group. The protocol therefore treats ledger adoption as conditional and permits a signed-registry profile for lower-complexity deployments [66], [68].

The proposal has not yet been implemented or evaluated with engineering learners, instructors, registrars, employers, or wallet vendors. It does not establish the optimal granularity of a micro-credential, identify a universally acceptable recency period, or prescribe how many institutions must endorse a governance transaction. It also does not address every legal retention, data-protection, or cross-border transfer requirement that may apply in a particular jurisdiction. Participating institutions must review those requirements in their operating context and should not assume that pseudonymization or hashing automatically removes a record from applicable privacy obligations.

11.3. Research and development agenda

The next technical step is a reference implementation that uses synthetic learners and independently implemented issuer, wallet, verifier, and registry components. It should test conformance, key compromise, selective disclosure, outage behavior, and status correlation under realistic traffic patterns. A separate educational-design study should co-create outcome maps for one engineering pathway and examine whether faculty reviewers agree

© 2026 TRTEE. Open Access under CC BY 4.0

on evidence sufficiency and whether learners understand the presentation and appeal workflow. A deployment evaluation should then compare administrative effort and decision traceability with existing transcript or badge processes without treating faster processing as proof of better learning [70], [72].

Further work should examine: privacy-preserving duplicate-evidence detection; accessible recovery for lost credentials; interoperability between achievement credentials and institutional records; governance models for consortium exit and long-term preservation; and whether shared registries can support interoperable mappings without encouraging a single universal competency taxonomy. Research should also test how pathway explanations affect learner agency and faculty trust. These questions are empirical and governance-sensitive [2], [16], [36]; the present article defines a design space and validation plan rather than claiming that a protocol has already answered them.

12. Conclusion

This article has specified a permissioned-blockchain and verifiable-credential protocol for stackable engineering micro-credentials. The design uses the ledger for consortium trust, version commitments, and auditable governance; it keeps learner claims, assessment records, and evidence off-chain; and it evaluates stackability through explicit outcome, assessment, prerequisite, recency, equivalence, and non-duplication rules. Cryptographic validity is necessary but not sufficient: the issuer must assess responsibly, the pathway steward must publish defensible rules, and the receiving institution must retain authority over credit and progression decisions. A synthetic laboratory-readiness scenario shows how claims from multiple providers can be verified and combined without creating a public learner transcript. Next comes a reference implementation and synthetic conformance testing. A limited pilot requires privacy, accessibility, and readiness review before real use. Deployment performance and educational effects remain unvalidated.

© 2026 TRTEE. Open Access under CC BY 4.0

Acknowledgements

The author thanks colleagues in the Faculty of Computer Science and Information Technology, Universiti Putra Malaysia, for discussion of an earlier outline of this article, and two anonymous reviewers whose comments substantially improved the treatment of governance in Section 7. No external funding was received for this article.

Author Contributions

Y.L.: Conceptualization, Methodology, Literature Synthesis, Writing — Original Draft, Supervision, Validation, Visualization, Writing — Review and Editing.

Conflicts of Interest

The author declares no conflicts of interest.

Generative AI Statement

Generative artificial intelligence was used for drafting and language revision.

© 2026 TRTEE. Open Access under CC BY 4.0

References

[1] S. Varadarajan, J. H. L. Koh, and B. K. Daniel, “A systematic review of the opportunities and challenges of micro-credentials for multiple stakeholders: learners, employers, higher education institutions and government,” International Journal of Educational Technology in Higher Education, vol. 20, no. 1, Art. no. 13, 2023. DOI: https://doi.org/10.1186/s41239-023-00381-x.

[2] N. T. N. Ha, M. Spittle, A. Watt, and N. Van Dyke, “A systematic literature review of micro-credentials in higher education: a non-zero-sum game,” Higher Education Research & Development, vol. 42, no. 6, pp. 1527–1548, 2023. DOI: https://doi.org/10.1080/07294360.2022.2146061.

[3] J. W. Prados, G. D. Peterson, and L. R. Lattuca, “Quality Assurance of Engineering Education through Accreditation: The Impact of Engineering Criteria 2000 and Its Global Influence,” Journal of Engineering Education, vol. 94, no. 1, pp. 165–184, 2005. DOI: https://doi.org/10.1002/j.2168-9830.2005.tb00836.x.

[4] W. M. Jones, S. Hope, and B. Adams, “Teachers' perceptions of digital badges as recognition of professional development,” British Journal of Educational Technology, vol. 49, no. 3, pp. 427–438, 2018. DOI: https://doi.org/10.1111/bjet.12557.

[5] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” Communications of the ACM, vol. 21, no. 2, pp. 120–126, 1978. DOI: https://doi.org/10.1145/359340.359342.

[6] H. A. Alsobhi, R. A. Alakhtar, A. Ubaid, O. K. Hussain, and F. K. Hussain, “Blockchain-based micro-credentialing system in higher education institutions: Systematic literature review,” Knowledge-Based Systems, vol. 265, Art. no. 110238, 2023. DOI: https://doi.org/10.1016/j.knosys.2022.110238.

[7] J. Sedlmeir, R. Smethurst, A. Rieger, and G. Fridgen, “Digital Identities and Verifiable Credentials,” Business & Information Systems Engineering, vol. 63, no. 5, pp. 603–613, 2021. DOI: https://doi.org/10.1007/s12599-021-00722-y.

[8] L. J. Shuman, M. Besterfield-Sacre, and J. McGourty, “The ABET “Professional Skills” - Can They Be Taught? Can They Be Assessed?,” Journal of Engineering Education, vol. 94, no. 1, pp. 41–55, 2005. DOI: https://doi.org/10.1002/j.2168-9830.2005.tb00828.x.

[9] B. M. Olds, B. M. Moskal, and R. L. Miller, “Assessment in Engineering Education: Evolution, Approaches and Future Collaborations,” Journal of Engineering Education, vol. 94, no. 1, pp. 13–25, 2005. DOI: https://doi.org/10.1002/j.2168-9830.2005.tb00826.x.

[10] S. Abramovich, C. Schunn, and R. M. Higashi, “Are badges useful in education?: it depends upon the type of badge and expertise of learner,” Educational Technology Research and Development, vol. 61, no. 2, pp. 217–232, 2013. DOI: https://doi.org/10.1007/s11423-013-9289-2.

[11] S. Roy, and D. Clark, “Digital badges, do they live up to the hype?,” British Journal of Educational Technology, vol. 50, no. 5, pp. 2619–2636, 2019. DOI: https://doi.org/10.1111/bjet.12709.

[12] T. J. Newby, and Z. Cheng, “Instructional digital badges: effective learning tools,” Educational Technology Research and Development, vol. 68, no. 3, pp. 1053–1067, 2020. DOI: https://doi.org/10.1007/s11423-019-09719-7.

[13] J. Biggs, “Enhancing teaching through constructive alignment,” Higher Education, vol. 32, no. 3, pp. 347–364, 1996. DOI: https://doi.org/10.1007/bf00138871.

[14] D. Boud, and N. Falchikov, “Aligning assessment with long-term learning,” Assessment & Evaluation in Higher Education, vol. 31, no. 4, pp. 399–413, 2006. DOI: https://doi.org/10.1080/02602930600679050.

[15] M. Risius, and K. Spohrer, “A Blockchain Research Framework,” Business & Information Systems Engineering, vol. 59, no. 6, pp. 385–409, 2017. DOI: https://doi.org/10.1007/s12599-017-0506-0.

[16] R. Beck, C. Müller-Bloch, and J. L. King, “Governance in the Blockchain Economy: A Framework and Research Agenda,” Journal of the Association for Information Systems, vol. 19, no. 10, pp. 1020–1034, 2018. DOI: https://doi.org/10.17705/1jais.00518.

[17] A. R. Hevner, S. T. March, J. Park, and S. Ram, “Design Science in Information Systems Research,” MIS Quarterly, vol. 28, no. 1, pp. 75–106, 2004. DOI: https://doi.org/10.2307/25148625.

[18] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, “A Design Science Research Methodology for Information Systems Research,” Journal of Management Information Systems, vol. 24, no. 3, pp. 45–77, 2007. DOI: https://doi.org/10.2753/mis0742-1222240302.

[19] J. T. M. Gulikers, T. J. Bastiaens, and P. A. Kirschner, “A five-dimensional framework for authentic assessment,” Educational Technology Research and Development, vol. 52, no. 3, pp. 67–86, 2004. DOI: https://doi.org/10.1007/bf02504676.

[20] P. Dawson, “Assessment rubrics: towards clearer and more replicable design, research and practice,” Assessment & Evaluation in Higher Education, vol. 42, no. 3, pp. 347–360, 2017. DOI: https://doi.org/10.1080/02602938.2015.1111294.

[21] L. D. Feisel, and A. J. Rosa, “The Role of the Laboratory in Undergraduate Engineering Education,” Journal of Engineering Education, vol. 94, no. 1, pp. 121–130, 2005. DOI: https://doi.org/10.1002/j.2168-9830.2005.tb00833.x.

© 2026 TRTEE. Open Access under CC BY 4.0

[22] J. Ma, and J. V. Nickerson, “Hands-on, simulated, and remote laboratories,” ACM Computing Surveys, vol. 38, no. 3, Art. no. 7, 2006. DOI: https://doi.org/10.1145/1132960.1132961.

[23] T. de Jong, M. C. Linn, and Z. C. Zacharia, “Physical and Virtual Laboratories in Science and Engineering Education,” Science, vol. 340, no. 6130, pp. 305–308, 2013. DOI: https://doi.org/10.1126/science.1230579.

[24] Y. Xiao, N. Zhang, W. Lou, and Y. T. Hou, “A Survey of Distributed Consensus Protocols for Blockchain Networks,” IEEE Communications Surveys & Tutorials, vol. 22, no. 2, pp. 1432–1465, 2020. DOI: https://doi.org/10.1109/comst.2020.2969706.

[25] J. Xu, C. Wang, and X. Jia, “A Survey of Blockchain Consensus Protocols,” ACM Computing Surveys, vol. 55, no. 13s, pp. 1–35, 2023. DOI: https://doi.org/10.1145/3579845.

[26] M. Castro, and B. Liskov, “Practical byzantine fault tolerance and proactive recovery,” ACM Transactions on Computer Systems, vol. 20, no. 4, pp. 398–461, 2002. DOI: https://doi.org/10.1145/571637.571640.

[27] J. H. Saltzer, and M. D. Schroeder, “The protection of information in computer systems,” Proceedings of the IEEE, vol. 63, no. 9, pp. 1278–1308, 1975. DOI: https://doi.org/10.1109/proc.1975.9939.

[28] D. F. Ferraiolo, R. Sandhu, S. Gavrila, D. R. Kuhn, and R. Chandramouli, “Proposed NIST standard for role-based access control,” ACM Transactions on Information and System Security, vol. 4, no. 3, pp. 224–274, 2001. DOI: https://doi.org/10.1145/501978.501980.

[29] S. Haber, and W. S. Stornetta, “How to time-stamp a digital document,” Journal of Cryptology, vol. 3, no. 2, pp. 99–111, 1991. DOI: https://doi.org/10.1007/bf00196791.

[30] Y. A. de Montjoye, L. Radaelli, V. K. Singh, and A. Pentland, “Unique in the shopping mall: On the reidentifiability of credit card metadata,” Science, vol. 347, no. 6221, pp. 536–539, 2015. DOI: https://doi.org/10.1126/science.1256297.

[31] M. Conti, E. Sandeep Kumar, C. Lal, and S. Ruj, “A Survey on Security and Privacy Issues of Bitcoin,” IEEE Communications Surveys & Tutorials, vol. 20, no. 4, pp. 3416–3452, 2018. DOI: https://doi.org/10.1109/comst.2018.2842460.

[32] A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, “A survey on essential components of a self-sovereign identity,” Computer Science Review, vol. 30, pp. 80–86, 2018. DOI: https://doi.org/10.1016/j.cosrev.2018.10.002.

[33] P. Dunphy, and F. A. P. Petitcolas, “A First Look at Identity Management Schemes on the Blockchain,” IEEE Security & Privacy, vol. 16, no. 4, pp. 20–29, 2018. DOI: https://doi.org/10.1109/msp.2018.3111247.

[34] S. Wang, L. Ouyang, Y. Yuan, X. Ni, X. Han, and F. Y. Wang, “Blockchain-Enabled Smart Contracts: Architecture, Applications, and Future Trends,” IEEE Transactions on Systems, Man, and Cybernetics: Systems, vol. 49, no. 11, pp. 2266–2277, 2019. DOI: https://doi.org/10.1109/tsmc.2019.2895123.

[35] C. Mazzocca, A. Acar, S. Uluagac, R. Montanari, P. Bellavista, and M. Conti, “A Survey on Decentralized Identifiers and Verifiable Credentials,” IEEE Communications Surveys & Tutorials, vol. 27, no. 6, pp. 3641–3671, 2025. DOI: https://doi.org/10.1109/comst.2025.3543197.

[36] Z. Yan, X. Zhao, Y. A. Liu, and X. R. Luo, “Blockchain-driven decentralized identity management: An interdisciplinary review and research agenda,” Information & Management, vol. 61, no. 7, Art. no. 104026, 2024. DOI: https://doi.org/10.1016/j.im.2024.104026.

[37] J. Camenisch, and T. Groß, “Efficient Attributes for Anonymous Credentials,” ACM Transactions on Information and System Security, vol. 15, no. 1, pp. 1–30, 2012. DOI: https://doi.org/10.1145/2133375.2133379.

[38] G. Fuchsbauer, C. Hanser, and D. Slamanig, “Structure-Preserving Signatures on Equivalence Classes and Constant-Size Anonymous Credentials,” Journal of Cryptology, vol. 32, no. 2, pp. 498–546, 2019. DOI: https://doi.org/10.1007/s00145-018-9281-4.

[39] A. Jonsson, and G. Svingby, “The use of scoring rubrics: Reliability, validity and educational consequences,” Educational Research Review, vol. 2, no. 2, pp. 130–144, 2007. DOI: https://doi.org/10.1016/j.edurev.2007.05.002.

[40] E. Panadero, and A. Jonsson, “The use of scoring rubrics for formative assessment purposes revisited: A review,” Educational Research Review, vol. 9, pp. 129–144, 2013. DOI: https://doi.org/10.1016/j.edurev.2013.01.002.

[41] D. R. Sadler, “Formative assessment and the design of instructional systems,” Instructional Science, vol. 18, no. 2, pp. 119–144, 1989. DOI: https://doi.org/10.1007/bf00117714.

[42] P. Black, and D. Wiliam, “Assessment and Classroom Learning,” Assessment in Education: Principles, Policy & Practice, vol. 5, no. 1, pp. 7–74, 1998. DOI: https://doi.org/10.1080/0969595980050102.

[43] D. J. Nicol, and D. Macfarlane-Dick, “Formative assessment and self-regulated learning: a model and seven principles of good feedback practice,” Studies in Higher Education, vol. 31, no. 2, pp. 199–218, 2006. DOI: https://doi.org/10.1080/03075070600572090.

© 2026 TRTEE. Open Access under CC BY 4.0

[44] J. Glöckler, J. Sedlmeir, M. Frank, and G. Fridgen, “A Systematic Review of Identity and Access Management Requirements in Enterprises and Potential Contributions of Self-Sovereign Identity,” Business & Information Systems Engineering, vol. 66, no. 4, pp. 421–440, 2024. DOI: https://doi.org/10.1007/s12599-023-00830-x.

[45] B. Libert, S. Ling, K. Nguyen, and H. Wang, “Zero-Knowledge Arguments for Lattice-Based Accumulators: Logarithmic-Size Ring Signatures and Group Signatures Without Trapdoors,” Journal of Cryptology, vol. 36, no. 3, Art. no. 23, 2023. DOI: https://doi.org/10.1007/s00145-023-09470-6.

[46] J. Tai, R. Ajjawi, D. Boud, P. Dawson, and E. Panadero, “Developing evaluative judgement: enabling students to make decisions about the quality of work,” Higher Education, vol. 76, no. 3, pp. 467–481, 2018. DOI: https://doi.org/10.1007/s10734-017-0220-3.

[47] D. P. Sidhu, and T. K. Leung, “Formal methods for protocol testing: a detailed study,” IEEE Transactions on Software Engineering, vol. 15, no. 4, pp. 413–426, 1989. DOI: https://doi.org/10.1109/32.16602.

[48] D. Boud, and E. Molloy, “Rethinking models of feedback for learning: the challenge of design,” Assessment & Evaluation in Higher Education, vol. 38, no. 6, pp. 698–712, 2013. DOI: https://doi.org/10.1080/02602938.2012.691462.

[49] D. Carless, and D. Boud, “The development of student feedback literacy: enabling uptake of feedback,” Assessment & Evaluation in Higher Education, vol. 43, no. 8, pp. 1315–1325, 2018. DOI: https://doi.org/10.1080/02602938.2018.1463354.

[50] X. Li, P. Jiang, T. Chen, X. Luo, and Q. Wen, “A survey on the security of blockchain systems,” Future Generation Computer Systems, vol. 107, pp. 841–853, 2020. DOI: https://doi.org/10.1016/j.future.2017.08.020.

[51] N. Ivanov, C. Li, Q. Yan, Z. Sun, Z. Cao, and X. Luo, “Security Threat Mitigation for Smart Contracts: A Comprehensive Survey,” ACM Computing Surveys, vol. 55, no. 14s, pp. 1–37, 2023. DOI: https://doi.org/10.1145/3593293.

[52] D. Chaum, “Security without identification: transaction systems to make big brother obsolete,” Communications of the ACM, vol. 28, no. 10, pp. 1030–1044, 1985. DOI: https://doi.org/10.1145/4372.4373.

[53] H. Nissenbaum, “A Contextual Approach to Privacy Online,” Daedalus, vol. 140, no. 4, pp. 32–48, 2011. DOI: https://doi.org/10.1162/daed_a_00113.

[54] S. Goldwasser, S. Micali, and C. Rackoff, “The Knowledge Complexity of Interactive Proof Systems,” SIAM Journal on Computing, vol. 18, no. 1, pp. 186–208, 1989. DOI: https://doi.org/10.1137/0218012.

[55] O. Goldreich, S. Micali, and A. Wigderson, “Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems,” Journal of the ACM, vol. 38, no. 3, pp. 690–728, 1991. DOI: https://doi.org/10.1145/116825.116852.

[56] O. Goldreich, and Y. Oren, “Definitions and properties of zero-knowledge proof systems,” Journal of Cryptology, vol. 7, no. 1, pp. 1–32, 1994. DOI: https://doi.org/10.1007/bf00195207.

[57] M. Prince, “Does Active Learning Work? A Review of the Research,” Journal of Engineering Education, vol. 93, no. 3, pp. 223–231, 2004. DOI: https://doi.org/10.1002/j.2168-9830.2004.tb00809.x.

[58] M. J. Prince, and R. M. Felder, “Inductive Teaching and Learning Methods: Definitions, Comparisons, and Research Bases,” Journal of Engineering Education, vol. 95, no. 2, pp. 123–138, 2006. DOI: https://doi.org/10.1002/j.2168-9830.2006.tb00884.x.

[59] S. Freeman et al., “Active learning increases student performance in science, engineering, and mathematics,” Proceedings of the National Academy of Sciences, vol. 111, no. 23, pp. 8410–8415, 2014. DOI: https://doi.org/10.1073/pnas.1319030111.

[60] V. J. Shute, “Focus on Formative Feedback,” Review of Educational Research, vol. 78, no. 1, pp. 153–189, 2008. DOI: https://doi.org/10.3102/0034654307313795.

[61] J. Hattie, and H. Timperley, “The Power of Feedback,” Review of Educational Research, vol. 77, no. 1, pp. 81–112, 2007. DOI: https://doi.org/10.3102/003465430298487.

[62] L. Wheelahan, and G. Moodie, “Analysing micro-credentials in higher education: a Bernsteinian analysis,” Journal of Curriculum Studies, vol. 53, no. 2, pp. 212–228, 2021. DOI: https://doi.org/10.1080/00220272.2021.1887358.

[63] L. Wheelahan, and G. Moodie, “Gig qualifications for the gig economy: micro-credentials and the ‘hungry mile’,” Higher Education, vol. 83, no. 6, pp. 1279–1295, 2022. DOI: https://doi.org/10.1007/s10734-021-00742-3.

[64] R. Desmarchelier, and L. J. Cary, “Toward just and equitable micro-credentials: an Australian perspective,” International Journal of Educational Technology in Higher Education, vol. 19, no. 1, Art. no. 25, 2022. DOI: https://doi.org/10.1186/s41239-022-00332-y.

[65] A. Acquisti, L. Brandimarte, and G. Loewenstein, “Privacy and human behavior in the age of information,” Science, vol. 347, no. 6221, pp. 509–514, 2015. DOI: https://doi.org/10.1126/science.aaa1465.

[66] F. Casino, T. K. Dasaklis, and C. Patsakis, “A systematic literature review of blockchain-based applications: Current status, classification and open issues,” Telematics and Informatics, vol. 36, pp. 55–81, 2019. DOI: https://doi.org/10.1016/j.tele.2018.11.006.

© 2026 TRTEE. Open Access under CC BY 4.0

[67] M. Janssen, V. Weerakkody, E. Ismagilova, U. Sivarajah, and Z. Irani, “A framework for analysing blockchain technology adoption: Integrating institutional, market and technical factors,” International Journal of Information Management, vol. 50, pp. 302–309, 2020. DOI: https://doi.org/10.1016/j.ijinfomgt.2019.08.012.

[68] E. Toufaily, T. Zalan, and S. B. Dhaou, “A framework of blockchain technology adoption: An investigation of challenges and expected value,” Information & Management, vol. 58, no. 3, Art. no. 103444, 2021. DOI: https://doi.org/10.1016/j.im.2021.103444.

[69] R. Belchior, A. Vasconcelos, S. Guerreiro, and M. Correia, “A Survey on Blockchain Interoperability: Past, Present, and Future Trends,” ACM Computing Surveys, vol. 54, no. 8, pp. 1–41, 2022. DOI: https://doi.org/10.1145/3471140.

[70] J. W. M. Lai, and M. Bower, “How is the use of technology in education evaluated? A systematic review,” Computers & Education, vol. 133, pp. 27–42, 2019. DOI: https://doi.org/10.1016/j.compedu.2019.01.010.

[71] J. Venable, J. Pries-Heje, and R. Baskerville, “FEDS: a Framework for Evaluation in Design Science Research,” European Journal of Information Systems, vol. 25, no. 1, pp. 77–89, 2016. DOI: https://doi.org/10.1057/ejis.2014.36.

[72] L. Deslauriers, L. S. McCarty, K. Miller, K. Callaghan, and G. Kestin, “Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom,” Proceedings of the National Academy of Sciences, vol. 116, no. 39, pp. 19251–19257, 2019. DOI: https://doi.org/10.1073/pnas.1821936116.

[73] E. J. Theobald et al., “Active learning narrows achievement gaps for underrepresented students in undergraduate science, technology, engineering, and math,” Proceedings of the National Academy of Sciences, vol. 117, no. 12, pp. 6476–6483, 2020. DOI: https://doi.org/10.1073/pnas.1916903117.

[74] Council of the European Union, “Council Recommendation of 16 June 2022 on a European approach to micro-credentials for lifelong learning and employability,” 2022/C 243/02, Official Journal of the European Union, C 243, pp. 10–25, 27 June 2022. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=oj%3AJOC_2022_243_R_0002

[75] W3C, “Verifiable Credentials Data Model v2.0,” W3C Recommendation, 15 May 2025. https://www.w3.org/TR/2025/REC-vc-data-model-2.0-20250515/

[76] W3C, “Decentralized Identifiers (DIDs) v1.0,” W3C Recommendation, 19 July 2022. https://www.w3.org/TR/2022/REC-did-core-20220719/

[77] 1EdTech Consortium, “Open Badges Specification v3.0,” version history, including document revision 1.4.1, 6 November 2025. https://www.imsglobal.org/spec/ob/v3p0

[78] W3C, “Bitstring Status List v1.0,” W3C Recommendation, 15 May 2025. https://www.w3.org/TR/2025/REC-vc-bitstring-status-list-20250515/

[79] W3C, “Verifiable Credential Data Integrity 1.0,” W3C Recommendation, 15 May 2025. https://www.w3.org/TR/2025/REC-vc-data-integrity-20250515/

[80] W3C, “Securing Verifiable Credentials using JOSE and COSE,” W3C Recommendation, 15 May 2025. https://www.w3.org/TR/2025/REC-vc-jose-cose-20250515/

[81] OpenID Foundation, “OpenID for Verifiable Presentations 1.0,” Final, 9 July 2025. https://openid.net/specs/openid-4-verifiable-presentations-1_0-final.html

[82] Hyperledger Fabric, “Endorsement policies,” release 2.5 documentation. https://hyperledger-fabric.readthedocs.io/en/release-2.5/endorsement-policies.html

© 2026 TRTEE. Open Access under CC BY 4.0