Community Trust Infrastructure

An Open Reference Architecture for Portable Operational Trust

Public research edition | Version 2.2 | November 2026

Open reference architecture • Working research manuscript for public critique and implementation

Financial infrastructure is becoming programmable and interoperable. Can the evidence and operational trust upon which decisions depend become portable and interoperable too?

Publication note: This public edition presents an evolving research architecture. Pilot observations are drawn from a team-review draft; operational configurations, participant-level records and sensitive details should not be inferred from this publication. Technical and program status may have changed since the source date.

Abstract

Community Trust Infrastructure (CTI) proposes an open, technology-agnostic reference architecture through which evidence of real-world activity can be captured, interpreted, validated and expressed as portable trust context. The problem is not simply that resources move too slowly. Communities can be active, organized and valuable yet remain difficult for outside actors to recognize. Existing institutions often rely on retrospective narratives, organizational credentials and repeated due diligence. CTI asks whether trustworthy evidence can emerge more continuously from ordinary coordination, at a lower burden to communities, while retaining local governance authority and selective disclosure.

This manuscript distinguishes record integrity from the meaning of a record; evidentiary validity from governance authorization; and trust infrastructure from the financial systems that allocate, represent and settle value. It treats the AIFS–Sarreya Tech365 pilot as an early reference implementation and critical source of negative as well as positive learning. SportFi and the Bolivia deployment are framed as prospective tests of funder confidence and portability, not completed proof. A separate Canada–EU Monetary Interoperability research branch is preserved as an exploratory macroeconomic hypothesis, outside the CTI reference architecture.

Research status and scope

CTI is a proposed reference architecture, not a deployed universal standard, a protocol, a blockchain, a scoring system, or a claim that trust can be objectively computed. Some

observations in this manuscript come from the AIFS × Sarreya internal pilot draft (Rajani and Lanigan, 26 August 2026); others are design propositions or hypotheses requiring comparative testing. The manuscript explicitly identifies these distinctions.

Part I — The architecture

1. The emerging trust constraint

Payments, identity, data exchange and programmable finance can make transactions easier without establishing why a particular claim, recipient or action deserves confidence. The central constraint is therefore not universally the speed of capital movement; in some settings it is recognition, legitimate authority, contextual evidence or access to the financial system. These constraints interact but should not be conflated.

The missing infrastructure in philanthropy may not be another mechanism for measuring impact after capital has been deployed, but a credible way of recognizing value communities are already producing before capital arrives.

This is a research proposition, not a finding about all philanthropy. It asks whether capital allocation can be informed by credible evidence of existing coordination without converting community life into a continuous external performance audit.

2. Definition and objective

Community Trust Infrastructure is an open reference architecture for portable operational trust: the ability to carry forward a calibrated, evidence-based understanding of how an actor, project or institution operates across contexts without requiring each new relationship to restart from zero. CTI standardizes functions and interfaces conceptually, not specific tools or universal judgments.

The intended portable output is trust context: a structured account of a claim, its supporting evidence, provenance, validation, relevant authority, limits, and conditions of use. A receiving actor remains responsible for interpreting that context against its own decision criteria. Portability of evidence does not guarantee identical conclusions.

3. Core research questions

- Can communities generate credible operational evidence as a low-burden byproduct of ordinary activity?
- What validation and authorization conditions allow that evidence to be evaluated outside its originating context?
- Can portable trust context reduce repeated proof costs while preserving privacy, autonomy and contestability?
- Does access to community-generated evidence measurably change funder confidence or capital-allocation decisions?

- Which architectural functions generalize across communities, technologies and jurisdictions, and which remain relational or context-specific?

4. Canonical functional sequence

Community activity → signals → attestations → validation → trust context → decision / policy interface → coordination or resource allocation

This is a logical sequence, not a mandatory software pipeline. Signals are observations or records; attestations are claims made by identifiable or contextually accountable parties; validation assesses both the evidence and the authority to make or confirm the claim; trust context preserves relevant qualifications. Decisions may concern participation, stewardship, access, governance, support or capital. Execution remains the responsibility of the relevant operational or financial system.

5. Verifiability, meaning and authority

Verifiability establishes that a record exists and has not been altered. It does not establish that the record means what we think it means.

More precisely, cryptographic verification can establish integrity or signer-related properties under particular assumptions; it cannot independently establish that an observed check-in was genuine participation, that participation produced a particular outcome, or that a signer had legitimate community authority. CTI therefore distinguishes the following questions:

Evidentiary validity and governance validity are independent. A claim may be factually correct but issued without authority; an authorized actor may issue a weak or incorrect claim. A trustworthy system should not depend on participants correctly identifying trustworthy people from charisma, familiarity, credentials or appearance. It should make relevant claims and authority assessable through accountable processes.

6. Architectural principles

- Reality first; signals should relate to actual activity, not only retrospective assertions.
- Signal before insight; capture simple, meaningful observations before interpreting them.
- Validation over assumption; technical existence and social reputation are insufficient alone.
- Distributed trust with legitimate authority; multiple signatures do not automatically establish legitimacy.
- Selective transparency; disclose what is necessary, protect what is sensitive.
- AI as interpreter, not authority; agents assist but do not acquire governance rights by technical capability.
- Low friction, high integrity; optimize sufficient confidence at the lowest reasonable proof cost.

- Local autonomy; communities define what counts, what is recorded or withheld, and who may access it.
- Resilience through human grounding; digital infrastructure supports but does not originate community relationships.
- Trust as maintained infrastructure; stewardship, corrections, appeals and long-term continuity require resources.

7. Architectural rules and governance

Decentralized governance does not require decentralized execution. It requires explicit boundaries around where authority resides, how it is delegated, what it permits, how actions are audited, and how authority can be modified or revoked. Human teams and agents may execute within bounded mandates while communities retain policy and governance authority.

Validation should carry an authority context: who appointed the validator, the scope and duration of the mandate, any permitted subdelegation, and its revocation status. A signature proves control of signing capability, not community authorization. CTI implementations should make the delegation path legible without imposing a single governance model.

Every major function should be replaceable; offline and low-tech workflows should remain possible; errors must be correctable or contestable; privacy and safety take precedence over maximal publication; and no single vendor or chain should be necessary to implement CTI.

Part II — Components and interfaces

8. Evidence and signal capture

Capture should occur as close to activity as feasible, using tools communities already use. A confirmed signal captured when an activity occurs can carry more evidential value, at lower community cost, than a retrospective narrative report. This is a conditional hypothesis: contemporaneous signals can also be misleading, incomplete or coerced. Narratives remain important for interpretation, explanation and learning, but need not carry the entire burden of proving that activity happened.

9. Attestation, validation and confidence

Attestation expresses a claim; validation tests the claim and the legitimacy of the process that produced it. Validation may involve peer or steward confirmation, supporting artifacts, checks for inconsistency, or independent review. Assurance should be proportionate to decision risk. The aim is sufficient confidence, not impossible certainty or a universal reputation score.

10. Trust context, portability and privacy

A portable trust context should identify the claim, evidence references, validation procedure, governance provenance, confidence limitations, disclosure permissions, and means of contesting or updating the record. It should support selective disclosure and minimize unnecessary personal information. Revoking an attestation is not equivalent to deleting public data, a distinction particularly important for images and information about women, children and vulnerable groups.

11. Human stewardship and agentic systems

AI can assist with organizing signals, flagging anomalies, summarizing evidence and preparing review queues. It should not silently substitute its own interpretation for community confirmation. At small scale, manual stewardship may be safer, cheaper and more responsive than automation. An agent's execution rights must be derived from explicit governance delegation and enforced in running code, not merely stated in documentation or prompts.

12. Coordination primitives and implementations

Prueba is an example attestation and record-structure primitive; Witness is a community-facing signal-capture workflow; POIDH may be investigated as a contribution-recognition primitive; identity tools may establish participation or authority; and settlement rails move value. None is CTI itself. Tokens, if used, should follow demonstrated contribution and legitimate community choices rather than create speculative incentives that distort participation.

13. The financial infrastructure boundary

CTI provides trust context that may inform decisions. It does not itself issue money, hold reserves or settle transactions. Capital allocation is the interface between trust/decision infrastructure and financial infrastructure. The distinct financial functions are:

These are functional distinctions, not a claim that all monetary systems use the same reserve structure or that assets always sit in a strict technical hierarchy. A conventional bank transfer and a stablecoin transfer may consume similar trust context while relying on different monetary and legal arrangements.

Part III — Reference implementations and evidence

14. AIFS × Sarreya Tech365: study setting

The internal draft "What we learned and what happens next" (Rajani and Lanigan, 26 August 2026) reports a USD 30,000 Olympism365 Tech365 pilot via Women Win, running September 2025–July 2026 with Sarreya Sport in Hargeisa, Somaliland. The project tested whether a community could create verifiable records of its activity and whether funders could treat those records as meaningful evidence. It included Prueba, a WhatsApp-based

Somali/English Witness agent, on-chain attestations, an indexer/dashboard, a USDC treasury and an off-ramp relationship.

The draft reports a May 15–16, 2026 workshop with 31 women, community agenda-setting, first-aid signups and a paper-ballot decision about USD 1,000 in project spending. These are documented project observations, not an independent impact evaluation.

15. What the implementation taught—and complicated

- Recognition before speed: the original payments-speed framing proved incomplete. Financial exclusion and legitimacy were major obstacles; better rails alone did not solve them.
- Confirmation is socially consequential: community agreement mattered more to participants than the choice of blockchain.
- Relational infrastructure matters: access to the ELAV Capital off-ramp came through personal introductions, exposing concentration and continuity risks.
- Manual stewardship was adequate at pilot scale: automation would not necessarily have improved the key workflows.
- Technical validity is insufficient: the v1 schema truncated IPFS identifiers, leaving some on-chain references unusable.
- Declared governance is not enforced governance: the pilot identified gaps between intended confirmation rules and deployed signing authority.
- Irreversibility creates operational risk: an address-poisoning incident exposed irreversible-transfer risk; public IPFS photographs pose lasting privacy risks despite revocable attestations.
- Documentation drift is a trust failure: inaccurate descriptions of messaging channels and revocation behaviour undermine reuse and auditability.

These are not peripheral implementation bugs. They directly challenge simplistic versions of the CTI thesis. A verifiable record may be broken, unauthorized, inaccurately described or unsafe to publish. The implementation therefore supplies design constraints and failure modes, not proof that CTI has already solved them.

16. Unfinished commitments and implementation safeguards

The internal pilot review identified priority safeguards: independently enforced confirmation rules, legitimate community-controlled signing authority, private handling of participant imagery and consent, schema and documentation consistency, and continuity plans. Implementation status must be verified before any production-readiness claim.

17. SportFi and Bolivia: tests still to be conducted

The internal draft describes SportFi as a Women Win-supported pilot planned through January 2027, with USD 10,000 for build work and USD 50,000 for a yield fund. Its primary research question is whether validated community activity changes funder confidence

and resource allocation, not whether a new financial instrument can be invented. Ayni Skate Project in Bolivia offers a contrasting context for testing portability across communication practices, stewardship structures and financial access.

Evaluation should define failure criteria before outcomes are known: evidence production burden; validator legitimacy; rate of erroneous or contested attestations; community control of signing and disclosure; funder decision changes attributable to evidence; bespoke engineering time; maintenance costs; and resilience if a steward, provider or project operator exits. The draft proposes more than two weeks of bespoke engineering for Bolivia as one practical warning threshold, not a universal portability standard.

Part IV — Research agenda and future directions

18. Philanthropic recognition before impact measurement

Philanthropy frequently evaluates outcomes after capital is deployed. CTI explores whether communities can first make their existing coordination and contributions credibly recognizable. Recognition should not be confused with monetization or an exhaustive impact metric. The proposition must be tested against possible harms: surveillance, incentives to game signals, unequal digital literacy, and a new burden of constant self-documentation.

19. Contribution recognition and non-extractive tokenomics

A future research track could examine POIDH-like contribution records as a basis for recognition, responsibilities, access or community-governed resources. Tokenization is optional. The guiding sequence is contribution → evidence → authorized recognition → community-defined rights or support, rather than token reward → induced behaviour. Any economic mechanism should test whether it strengthens or distorts the activity it seeks to recognize.

20. Trust infrastructure as public-interest stewardship

Stewards, coordinators, validators, maintainers and community organizers generate the conditions for evidence to be credible. Their labour should be budgeted, not treated as invisible volunteer overhead. The central research challenge is whether proof costs fall for communities once the full cost of maintenance, privacy protection, appeals, training and governance is counted.

21. Resilience, institutional plurality and algorithm agility

CTI should permit heterogeneous tools, human-led fallback processes and migration of technical components as security conditions change. Cryptographic algorithm agility, data portability, revocation and continuity plans are implementation requirements to investigate. Neither blockchain permanence nor reliance on a single dashboard should be mistaken for durable community control.

22. Separate research branch: Canada–EU Monetary Interoperability

Status: exploratory macroeconomic research; not a CTI architectural requirement or demonstrated application.

Canada–EU monetary integration may not require monetary union. Digitally interoperable sovereign currencies, tokenized deposits and assets, common settlement infrastructure, and trusted verification systems could enable substantially deeper economic integration while preserving independent CAD and EUR monetary policy.

This hypothesis requires its own evidence concerning payment systems, cross-border settlement, foreign exchange, regulatory compatibility, financial stability, monetary sovereignty and actual adoption incentives. CAD and EUR could remain separate units of account and monetary-policy domains even where financial infrastructures become more interoperable. Such integration is not automatic: legal, liquidity, exchange-rate, supervision and governance barriers may be substantial.

CTI might someday supply a limited evidence/authorization interface to particular cross-border allocation or asset workflows. The monetary-interoperability thesis does not depend on CTI, and CTI should not be asked to prove it. This branch should be published and evaluated separately until a concrete interface and use case are established.

23. Open questions and falsification tests

- Does community-generated evidence actually alter funder decisions compared with conventional reporting?
- What minimum information is necessary to assess a claim across governance contexts without exposing private participants?
- Can authorization be portable while revocation and local autonomy remain effective?
- Can the model work where no stable steward pair or trusted relationship network already exists?
- Do savings in reporting effort exceed the total cost of verification and maintenance?
- How should disputed, false, obsolete or unsafe records be corrected and communicated downstream?
- When does contribution recognition create incentives that corrupt the signal?
- Can implementations survive loss of the original funder, indexer, off-ramp or technical maintainer?

24. Institutional comparison: World Bank FundsChain

The World Bank Group's FundsChain initiative is a useful institutional case for comparing transaction traceability and programmable financial controls with CTI's proposed evidence and authority functions. Publicly described implementation and evolution updates (2024–2026), including a reported Ethiopia pilot, indicate a progression from auditability and supporting-document verification toward purpose-bound financial flows.

These details require primary-source verification before being cited as independently established implementation results.

FundsChain and CTI address different, potentially complementary problems. Financial systems can document transactions and enforce conditions; CTI asks whether the real-world claims used to satisfy those conditions have sufficient provenance, contextual validity, legitimate authorization and community control. A cryptographically valid financial record does not itself establish that a community-defined outcome occurred.

The research gap is whether institutional systems can consume independently validated community activity attestations, community-defined verification parameters or distributed evidence governance without absorbing authority over the originating community. No such interoperability is assumed or claimed here.

Research implication: as capital becomes more programmable, evidence governance becomes more consequential. FundsChain should be treated as a high-priority institutional comparison, not a CTI reference implementation or proof of interoperability.

Conclusion

CTI is best understood as a proposed coordination architecture, not a promise that data, cryptography or automation can manufacture trust. Its contribution is to specify how real-world evidence, legitimate authority, contextual validation and portable trust context might help independent actors coordinate with sufficient confidence. The Sarreya pilot shows both why this question matters and why its solution cannot be presumed: relationships unlocked financial access; community confirmation carried legitimacy; and technical implementation gaps revealed how easily an apparently verifiable system can fall short of its own governance commitments.

The next research threshold is empirical. SportFi must test whether community-generated evidence changes decisions, while cross-context deployments test portability, cost and control. CTI should evolve through implementation, independent scrutiny and openly documented failure—not by absorbing every adjacent technological or macroeconomic hypothesis into its architecture.

Source and evidence register

Primary internal source: Rajani, H., and Lanigan, S. (26 August 2026). “What we learned and what happens next: AIFS × Sarreya, Tech365 pilot.” Draft v1 for team review (unpublished). Summarized here with attention to participant privacy and operational security; observations are not independently audited.

Prior project source: “Community Trust Infrastructure (CTI), V2.0 Compiled Manuscript” (working project manuscript). Used for continuity of problem statement, trust thesis, architecture and principles.

Research-development inputs: CTI project notes on portable operational trust, governance authorization, proof cost, selective transparency, financial infrastructure boundaries, contribution recognition and Canada–EU monetary interoperability. These are propositions requiring further sourcing and testing.

Institutional comparison to verify: World Bank Group (2024–2026), FundsChain implementation and evolution updates, including reports of FundsChain 2.0 testing in Ethiopia. Source classification: Institutional Case Studies — Programmable Development Finance and Verification Infrastructure. The specific public source URLs and claims should be checked before scholarly citation.

Evidence gap: this edition does not claim a completed literature review, external validation of pilot outcomes, measured cost savings, or demonstrated changes in funder allocation behaviour. A subsequent scholarly edition should attach the CTI Literature Index, primary technical artifacts, formal methods and independent sources to individual claims.