

Make abundance profitable.

An economy with a heart.

A first-principles synthesis of the LOVE protocol and the \$EXIT seed: economics, infrastructure, human oversight, and cultural adoption.

October 8, 2026

Blueprint for review and testing. No currency, investment product, or pilot has launched through this publication.

The argument

- 01 / Capitalism, connected to a heart.
- 02 / What becomes scarce when capability gets cheap?
- 03 / A north star, with a humility clause.
- 04 / The missing bridge: outcomes to revenue.
- 05 / One economic language. Many human worlds.
- 06 / A receipt for making life better.
- 07 / The substrates beneath the promise.
- 08 / AI proposes. People remain accountable.
- 09 / The song opens the door. The receipt keeps it open.
- 10 / The first 90 days: a pilot people can inspect.
- 11 / Rewards do not refill themselves.
- 12 / Failure modes belong in the founding document.
- 13 / From a seed to an institution.
- 14 / Reference design: the smallest honest machine.

Original seed supplied by Jack. Economic and technical foundation: LOVE Complete Edition v2.0, The Next Bitcoin Institute, July 2026, draft for comment. AI-assisted synthesis and implementation.

Capitalism, connected to a heart.

Loveitalism is an economic design principle: make improving other people's lives a competitive way to earn, build, and own. Its north star is an economy in which the strongest businesses expand access to the essentials of a good life, and the people creating that value share in the upside. Love is the objective. Markets are one coordination tool. Proof is the discipline.

The seed behind \$EXIT says that capital chases the fastest-rising asset. The more precise statement is that investors seek expected returns subject to risk, liquidity, time horizon, and constraints. Some capital chooses safety over appreciation; some is mandated to pursue social goals. An abundance asset therefore has to win on credible economics as well as moral appeal.

The design task is to create a valuable system whose revenues and durable advantages grow when useful outcomes grow. Price appreciation is a possible market consequence, not the mechanism and not a promise. A token that rises while its community gets poorer would fail the thesis, however impressive its chart.

Here is the compact proposition: produce more useful capability; make access broader; verify the improvement; collect payment from someone who benefits; pay the people who delivered it; reinvest in the next improvement. Culture turns that operating loop into something people want to belong to.

What becomes scarce when capability gets cheap?

Start with people, not tokens. People want security, health, shelter, agency, relationships, and meaningful work. Resources, knowledge, institutions, and coordination help meet those wants. Technology can lower particular costs, but land, energy, attention, physical construction, ecological limits, and political power do not disappear because software improves.

As an input gets cheaper, some value can shift to complementary inputs that make it useful. Cheap intelligence could increase the importance of trusted judgment, distinctive taste, distribution, and coordinated demand. This is a conditional economic argument, not a theorem that virtue becomes exponentially valuable or that every industry's marginal cost goes to zero.

An AI can draft a hundred companies. It cannot infer a community's legitimate priorities from throughput alone. Making a company work also requires customers willing to pay, trustworthy execution, people willing to participate, and decisions under uncertainty. Human oversight is useful because the objective itself is negotiated, local, and alive.

The relevant scarcity is often credible coordination: who wants this, who will pay, who can deliver, and who can establish that the result was worth it? Loveitalism organizes those questions. AI expands the capacity to answer them; accountable people retain the authority to choose and challenge the answers.

A north star, with a humility clause.

The desired end state is broad access to essential goods, rising real purchasing power, productive ownership shared with contributors, and institutions that reward improvement over manufactured scarcity. Success means a household can live better with less compulsory effort, while builders can prosper by making that improvement possible.

Call this a directional north star rather than a guaranteed destination. Lower costs, unmet needs, new coordination tools, and demand for meaning create an opportunity. They do not prove that abundance wins every market, that one network will dominate, or that LOVE will become a global currency. Institutions can capture productivity gains, and dominant platforms can preserve rents.

The scorecard therefore measures access and lived outcomes before market capitalization: essential costs relative to household income, quality and durability of service, beneficiary agency, additional outcomes per dollar, contributor compensation, concentration of control, and continued performance after subsidies decline. Growth that worsens these measures is drift, not victory.

Layer Won names the moment competing incentives become more cooperative. It is the anthem's horizon, not a claim that history ends or conflict disappears. \$EXIT means exiting the incentive trap by changing what earns a return. In this blueprint, it is a cultural title; it is not a launched financial instrument.

The missing bridge: outcomes to revenue.

A socially valuable result is not automatically a profitable result. A free tutorial can help thousands of people and still earn its maker nothing. An energy retrofit can generate savings and still fail commercially if the customer cannot finance it or the installer cannot collect. The mechanism needs a named payer, a measurable benefit, a workable contract, and affordable verification.

Three business patterns can connect abundance to revenue. A performance contract shares measured savings, as with energy efficiency. A subscription or service contract pays for reliable access and coordination, while a free common layer spreads knowledge. An institutional outcome contract pays for a defined result that a government, employer, foundation, or utility values. The payer and the beneficiary may be different people; those relationships need consent and explicit safeguards.

For a period, let collected revenue be R , contributor payouts W , verification V , operations O , and expected losses plus reserve funding L . Operating surplus is $M = R - W - V - O - L$. A viable recurring service needs M at least zero without endlessly increasing subsidy. Beneficiary value must also exceed the beneficiary's cost, including time, privacy loss, and imposed risk. Positive operator margins alone are insufficient.

Suppose a sponsor pays \$10,000 for a program that produces independently estimated benefits of \$16,000. An illustrative budget pays \$6,500 to contributors, \$500 for verification, \$2,000 for operations, and \$1,000 to reserves. Verification consumes 7.7% of contributor payouts; there is no unallocated operating surplus. These are scenario inputs, not pilot results. If measured benefit falls to \$7,000, the budget still balances but the payer's case fails. Accounting balance and social value are separate tests.

The asset connection must be specified independently. Equity may carry a legal claim on enterprise cash flow. A contribution agreement may compensate a builder for work. A currency may have transactional demand. A public impact record has informational value. None gives a token holder ownership of savings, revenue, or somebody else's good deed unless an enforceable agreement actually provides it. Financial participation belongs behind tested economics and the appropriate legal structure.

One economic language. Many human worlds.

The supplied LOVE Complete Edition v2.0 proposes a fixed-supply participation currency, an earning pool, a non-transferable reputation ledger called REP, community treasuries, and a social feed of verified deeds. Loveitalism takes that as a proposed implementation beneath a wider philosophy. The philosophy should survive even if its first currency design fails.

The source's genesis allocation of 8 billion LOVE, its percentage splits, fee rate, release caps, and reputation half-life are historical draft parameters. This new blueprint does not ratify or implement them. Allocation outside the earning pool must be disclosed plainly: reserving tokens for founders, liquidity, or treasury funding is different from earning them through verified work.

If later adopted, LOVE would be the transferable economic unit; REP would record demonstrated competence in a particular domain. Money balances and governance standing would have different purposes. Holding would not promise a yield. REP would be earned, non-transferable, appealable, and subject to published decay rules. A capable newcomer needs a path to earn standing, and a worker needs a way to take leave without losing the practical ability to participate.

Separating REP from money reduces direct purchase of voting weight. It does not make bribery, coercion, reputation laundering, favoritism, or collusion impossible. Domain limits, conflict disclosures, independent review, capped influence, transparent selection, and appeal routes remain necessary. Truth cannot be secured by declaring a second ledger virtuous.

Communities express identity through their names, feeds, rules, treasuries, and local contribution records. They do not need a separate speculative coin. Collateral-backed community shares are an optional later research path, not a launch requirement. Locked reserves can be redeemed; they are reversible demand, not a one-way supply sink or guaranteed source of appreciation. Liquidity, withdrawal rights, reserve accounting, and loss allocation must be explicit.

A receipt for making life better.

A useful claim has a full lifecycle: define, fund, deliver, review, challenge, settle, and monitor. The contract records a baseline, outcome standard, evidence requirements, payment rule, measurement period, consent, and dispute process before the work begins. It also names the payer, claimant, beneficiary, verifier, and the people authorized to release funds.

Evidence is a bundle. It can include invoices, signed observations, before-and-after measurements, provider records, and media whose provenance can be checked. A hash establishes that a file has not changed relative to the committed bytes; it does not establish that the event happened. EAS records attestations; it does not establish their truth. C2PA records provenance and tamper evidence; it does not settle the social value or factual validity of a scene.

AI helps organize evidence, detect duplicates, estimate uncertainty, and route review. A model suggestion never authorizes payment by itself. Low-value, low-risk claims may use automatic rule checks plus independent sampling. Large, contested, safety-sensitive, or suspicious claims require domain-qualified human review. Use a private evidence store with access controls; public receipts contain only consented, necessary information.

Challenges must be accessible to people with little money. A bond can deter spam, but it can also deter a truthful complainant. Subsidized appeals, protected reporting, independent reviewers, and proportionate limits matter as much as stake. Penalizing disagreement with the majority is not the same as penalizing dishonesty. Reviewers should be paid for competent review rather than agreement alone.

Payment follows acceptance and the required challenge period. A settled receipt identifies the standard, evidence summary, reviewer roles, date, payout, and any uncertainty or reversal. Later quality failures must be reportable. The system distinguishes completed action from additional impact: proving that work occurred does not prove it would not have happened anyway. Baselines, comparison groups where appropriate, and follow-up measurement improve that inference.

The substrates beneath the promise.

The physical substrate is useful production: reliable energy, housing, food systems, transport, health services, and education. A ledger cannot substitute for a licensed installer, a functional supply chain, or permission to act on somebody else's property. Start where delivery and measurement are tractable, then expand.

The economic substrate is funded demand: payer contracts, clear payment terms, risk allocation, an operating budget, and reserves. The identity substrate is accountable participation with plural routes and minimal disclosure. The proof substrate is versioned outcome standards, evidence custody, independent review, random audits, and appeals. The settlement substrate is conventional payment rails and escrow where available, with reconciliation and reversal procedures.

The governance substrate is explicit authority: who sets standards, who can change budgets, who has a conflict, how a decision is appealed, and what happens when the operator fails. The intelligence substrate is auditable AI assistance under human-set objectives. The cultural substrate is identity, music, story, rituals, and local conveners. These are all functional components; culture is how people learn which action is worth taking.

For the first pilot, use a mobile-friendly web app, a relational database with an append-only event log, encrypted object storage, passkeys and accessible recovery, a review queue, and a reconciled conventional payout provider. Keep personally identifying evidence off public chains. Build export and operator migration into the schema. This is cheaper to inspect and repair than launching a new chain before the operating model is understood.

Candidate open foundations include EAS for attestations, Open Source Observer for public impact-data patterns, cadCAD for mechanism simulation, and Colony for domain reputation and decay concepts. Safe can supply treasury-control patterns if an audited chain-based treasury is later justified. Review licenses, release tags, audits, integration cost, and project health before adoption. These are researched candidates, not integrations already running on this website.

Cryptographic agility is a real long-term requirement. NIST has standardized post-quantum algorithms, but adding one algorithm to account software does not make a whole chain, sequencer, bridge, custodian, or device post-quantum secure. Threat modeling, supported implementations, key rotation, recovery, and system-wide migration are separate engineering tasks. Do not market immunity or loss-proof recovery.

AI proposes. People remain accountable.

AI accelerates the search space; people make the commitment. The human roles are goal setting, taste, beneficiary representation, relationship building, interpretation, adjudication, and revision. These roles are valuable because their outputs must be accepted by actual people under actual constraints. Automation does not remove the need for legitimacy.

The product should preserve that distinction in every workflow. Show a model's confidence and supporting evidence. Record which model and rule version produced a suggestion. Require authorized human approval at consequential boundaries. Enable a beneficiary to contest an output. Compare model judgments with expert re-review, and test whether errors concentrate among particular groups or unfamiliar domains.

Culture also needs oversight. An AI can generate a hundred slogans, but a local organizer knows which one feels insulting, aspirational, or true. A creator knows when a chorus earns its emotion. A founder knows when a business claim overstates what exists. Those are operating skills, not decorative jobs. Fund them in the budget and measure whether they produce sustained participation and trust.

The goal is not a perfect autonomous company. It is a company that becomes more useful as its people learn, with automation helping them spend less effort on routine work. A constitution, logs, review rights, and accountable operators make that learning inspectable. The website presents the blueprint; it does not run an AI adjudicator or execute payments.

The song opens the door. The receipt keeps it open.

The cultural obstacle is not merely getting people to understand a mechanism. It is getting them to recognize themselves inside a possible future. A movement spreads when a person can perform a small act, feel a real result, tell the story honestly, and invite somebody else. Financial language alone rarely supplies that experience.

Loveitalism's promise is personal: your ambition can make somebody else's life better, and that can make yours better too. Its shared line is "More for you. More for me." Its verb is make. Its recurring artifact is a receipt for work that mattered. Its anthem is \$EXIT / Layer Won. The music makes the future feel possible; a local result makes it believable.

The first repeatable story is a 15-to-30-second settlement clip: here was the problem; here is what changed; here are the people who did it; here is the evidence link and the payment, if the claimant consents. Label submitted, under review, settled, and disputed claims clearly. Never portray a pending claim as a proven victory. Do not make people expose their hardship to qualify for help.

Start with one scene, not humanity at once: a campus, creator circle, neighborhood association, or existing builder community that already trusts a convener. Recruit the payer and delivery partners before attracting claimants. A weekly settlement gathering celebrates completed work, thanks reviewers, and tells the next mission. Small joyful projects teach the habit; a serious outcome contract tests the economics. Keep those two budgets distinguishable.

Creators are translators and conveners. Pay for their work through clear fees and consented deliverables, not undisclosed token incentives or recruitment commissions. Let musicians, designers, organizers, and beneficiaries adapt the language without turning participants into a moral leaderboard. Do not call the audience less caring because it has doubts. A movement worthy of the name makes dissent and exit possible.

Measure qualified participation rather than raw reach: viewers who open the evidence, eligible people who complete a first task, repeat doers, retained beneficiaries, matched funding converted into settled work, and costs per additional outcome. Songs, remixes, events, and documentary clips are proposed channels to test. None is a forecast of virality. The culture succeeds when it produces useful behavior that lasts after the launch excitement.

The first 90 days: a pilot people can inspect.

The proposed economic wedge is household energy savings delivered by qualified local providers. It offers a recognizable benefit, measurable usage, and a possible payer with an economic reason to participate. No pilot, sponsor, household, utility partner, or licensed provider has been secured by this publication. Those commitments are the first gate.

Days 1–30: choose one locality and one contractual payer. Recruit a community convener and qualified delivery partner. Define consent, baseline usage, weather and occupancy adjustments, eligibility, the savings window, and how rebound effects are treated. Pre-register success and stopping conditions. Use historical or parallel controls where feasible; ask an independent measurement specialist to review the design. Secure money for payouts, operations, audit, and the full follow-up period.

Days 31–60: onboard a small voluntary cohort, with a planning ceiling of 100 households rather than a recruitment promise. Offer understandable terms and a non-biometric identity route. Test the full claim lifecycle with conventional payments. Share only consented evidence. Publish operational measures and corrections weekly. No speculative token is required to enter, earn, verify, or leave.

Days 61–90: independently re-review a sample, evaluate estimated savings and beneficiary satisfaction, reconcile every payout, and publish a report with uncertainty. Energy results may require a longer or seasonal measurement period; day 90 is an operational review, not a promise to establish annual savings. Test a second cohort only when the contract, review cost, and trust measures support it.

Suggested planning gates, to be finalized before recruitment: verification costs no more than 10% of contributor payouts; all payouts reconciled; positive contribution margin under the actual payer contract; a qualified independent reviewer; zero unresolved severe consent or safety incidents; and at least 30% of eligible participants choosing another task or a follow-up invitation. These are proposed targets, not validated benchmarks.

Pause payouts on a severe integrity or safety failure. Redesign if verification remains above 25% of payout after two revisions, if the paying customer will not renew on demonstrated value, or if follow-up shows benefits disappear. Publish misses and uncertainty. Do not launch a token to finance a failing pilot. The next product is an institutional allocation service only if the evidence says the first service deserves to grow.

Rewards do not refill themselves.

A fixed supply cap limits the number of units created; it does not limit dilution from previously reserved units entering circulation. It also does not pay for operations. Distinguish total supply, circulating supply, pool balance, committed liabilities, and the operator's cash runway in every report.

For a reward pool, $P_{\text{next}} = P_{\text{now}} + \text{collected fee recycling} + \text{external funding} - \text{settled rewards} - \text{other authorized outflows}$. Non-negative balances today do not prove long-term sustainability. If rewards exceed inflows, runway is finite. If fees are collected in a volatile token while payroll is due in dollars, exchange risk creates an additional constraint. Simulate low adoption, adverse prices, provider outages, collusion, and a sudden withdrawal of subsidies.

Increasing transaction volume by encouraging wash activity can make fee dashboards look healthy while destroying economic value. Count economically useful demand, net revenue after rebates, beneficiary value, and outside payer renewal. A fee should pay for a useful service rather than function as a toll on every act of care. Publish who pays it and what they receive.

Any later token decision needs a specific necessity test: does it solve coordination or ownership problems better than conventional accounts and contracts, after integration cost, legal obligations, and risk? If the answer is no, do not add it. If yes, separate market price from operational health, stress-test redemptions where any exist, and specify rights honestly. No appreciation, passive yield, peg, or universal redemption is promised here.

Failure modes belong in the founding document.

Metric gaming: people optimize the paid proxy instead of the desired outcome. Use pre-registration, independent baselines, random audit, anti-duplication rules, and follow-up. Still report residual uncertainty. Manufactured need is fraud when the rules specify it; preventing it remains an operational challenge.

Collusion and capture: claimants, reviewers, insiders, or community leaders cooperate to redirect funds. Separate roles, disclose conflicts, rotate panels, re-review samples, cap influence, and maintain an appeal route outside the local group. An immutable record of a bad decision is still a bad decision.

Exclusion and surveillance: an identity gate can penalize people without documents, compatible devices, reliable connectivity, or safe public visibility. Provide multiple routes, assisted onboarding, minimal collection, retention limits, and private beneficiary evidence. Reputation must not become a universal rating of a person's moral worth.

Cultural drift: price hype, hero worship, coercive positivity, and public spectacle can replace useful work. Keep receipts near claims, disclose sponsorship, make disagreement welcome, and refuse promotion that implies guaranteed prosperity. The anthem's victory is a shared aim, not evidence that a market will reward a buyer.

Legal and contractual risk: payments, prizes, custody, transferable assets, prediction products, employment, taxes, and data protection involve different obligations. Classification depends on actual rights, operations, and jurisdiction. This paper sets design gates; it does not establish that a structure is lawful. Current launch-specific review and appropriate licensed partners are needed before those products operate. The public site offers reading, a scenario calculator, and lyrics, with no investment solicitation or transaction flow.

From a seed to an institution.

Stage A is this publication: a philosophy, engineering blueprint, cultural plan, and anthem. Stage B is a funded conventional-payment pilot with measurable outcomes and independent review. Stage C is a repeatable allocation product that paying institutions renew because it delivers better value. Stage D, only if justified, is an interoperable community network with shared standards and exportable histories. Stage E is an optional currency or shared-ownership architecture after the necessity test, economic evidence, security review, and relevant legal gates.

The core team needs mechanism design, product engineering, verification and measurement, community operations, cultural direction, and appropriate legal and accounting support. The creator-founder role is not a marketing accessory. It connects the economic thesis to a lived identity, tests emotional resonance, and keeps promises aligned with what the product can deliver. Jack's work belongs in that cultural direction and orchestration role.

The operating institution needs a funded payroll, clear service obligations, a published budget, and continuity beyond its founder. Open standards, exportable data, independent interfaces, and a practical migration path make exit credible. A community should be able to leave without losing its consented evidence and legitimate earned records, subject to privacy and retention requirements.

Loveitalism is not a finished answer to every problem in economics. It is a testable choice about what ambition should build: systems in which one person's progress increases the practical possibilities available to others. The machine can keep running. Its objective can change. Make abundance profitable, and prove it one contract, one outcome, one human story at a time.

Reference design: the smallest honest machine.

This reference architecture is a specification for a future pilot, not running protocol software. The authoritative record lives in a relational database. Every consequential change also creates an append-only event with `event_id`, `object_id`, `actor_id`, `timestamp`, `rule_version`, `previous_event_hash`, and `payload_hash`. Hash chaining helps reveal edits; signed external checkpoints and independent backups reduce the risk of an operator rewriting the entire history. Neither proves the original claim true.

An OutcomeContract contains `contract_id`, `payer_id`, `delivery_domain`, `budget_currency`, `funded_amount`, `maximum_liability`, `baseline_definition`, `outcome_metric`, `measurement_window`, `eligibility_rule`, `evidence_schema_version`, `payout_function`, `reviewer_policy`, `challenge_deadline_rule`, `consent_policy`, and `dispute_authority`. A Claim contains `claim_id`, `contract_id`, `claimant_id`, `beneficiary_consent_reference`, `baseline_reference`, `evidence_manifest_hash`, `requested_amount`, `status`, `created_at`, and `idempotency_key`. A Review records reviewer identity and domain qualification, conflicts, evidence references, finding, rationale, and rule version. A Payout records the accepted claim, authorized amount, provider reference, currency, and reconciliation state.

The allowed claim progression is draft to submitted to under_review to accepted to payable to paid. A submission can instead become rejected or disputed. Accepted means a finding has been made; payable means the required challenge period has ended and funding remains reserved. Disputed claims cannot become payable until the named independent authority resolves them. Later reversal is recorded as a new event and handled under the contract; history is not deleted or silently relabeled.

Enforce these invariants inside database transactions: outstanding authorized payouts never exceed funded available money; each claim has one active settlement version; each payment attempt has a unique idempotency key; a claimant cannot serve as their own independent reviewer; AI output cannot supply the required approval signature; a paid status requires the payout provider's confirmed result and ledger reconciliation; private evidence never appears in a public receipt without the required consent. Provider callbacks are authenticated, deduplicated, and processed through a retryable queue. A timeout produces pending investigation, not a second payment.

The minimum service interfaces are: create a funded contract; submit a claim and evidence manifest; propose or record a review; open a dispute; authorize a payable claim; reconcile a payout; and retrieve a redacted public receipt. Authentication and role checks apply to every write. Sensitive files use scoped, expiring access and malware checks; original evidence has retention and deletion rules. Audit-log exports omit unnecessary personal data. Disaster recovery is rehearsed with a restore test

and a documented secondary operator.

If a later LOVE earning pool is justified, the source draft's monthly release pattern can be tested as $\text{release} = \min(\text{governed_fraction times available_pool}, \text{eligible_claim_liabilities}, \text{funded_epoch_budget})$. Unpaid eligible claims remain queued obligations with clearly stated currency and payment terms; a pool cap must not silently convert promised compensation into indefinite waiting. Pool accounting and operational cash reserves remain separate. Proposed REP events are derived from accepted, audited contributions and can be corrected through appeals; they do not authorize unreviewed financial actions.

Acceptance tests for an implementation should exercise duplicate callbacks, a depleted budget, a late challenge, a conflicted reviewer, revoked consent, an unavailable provider, an appeal that changes the finding, and recovery from backup. Before real payments, prove that the system prevents double settlement and preserves every authorized liability. Before claiming impact, independently re-review the measurements. Before claiming cultural fit, observe people voluntarily returning to the work.

Sources & foundations

Primary sources checked during this edition. They establish components and precedents, not validation of this complete system. The source draft's numerical incident claims and legal conclusions are not repeated as verified current facts.

The supplied foundation

https://drive.google.com/file/d/1myTOXh7EED_O_VsS9zze0EHnpeUDNh4M/view

LOVE: A Merit-Backed Currency and the Architecture of the Socialverse, Complete Edition v2.0. The Next Bitcoin Institute, July 2026; draft for comment. Read in full for its protocol, principles, and cultural design. Its empirical and legal claims are not automatically ratified here.

Ethereum Attestation Service

<https://github.com/ethereum-attestation-service/eas-contracts>

Open schemas and attestation contracts; MIT. Attestation infrastructure, not a truth oracle.

Open Source Observer

<https://github.com/opensource-observer/oso>

Open data and impact-analysis infrastructure. A reference for reproducible measurement pipelines, not a complete physical-world verification product.

Optimism Retro Funding

<https://github.com/ethereum-optimism/Retro-Funding>

Published funding data and methods. A precedent for inspectable allocations, not proof that this new mechanism will work.

cadCAD

<https://github.com/cadCAD-org/cadCAD>

MIT simulation framework for complex systems. A candidate for parameter sweeps and adversarial economic models.

Colony reputation

<https://docs.colony.io/colonynetwork/tldr/reputation/>

Domain reputation and decay as design precedents. Governance standing still requires anti-capture protections.

Safe Smart Account

<https://github.com/safe-fndn/safe-smart-account>

LGPL-3.0 treasury and smart-account code. Use only appropriate reviewed releases if a chain-based treasury becomes necessary.

Human Passport

<https://github.com/passportxyz/passport>

An identity and anti-Sybil candidate. A credential alone is not sufficient proof of unique, honest participation.

C2PA explainer

<https://spec.c2pa.org/specifications/specifications/2.3/explainer/Explainer.html>

Primary guidance explicitly distinguishes media provenance from judgments about truth. Relevant to evidence handling.

NIST post-quantum standards

<https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>

Primary announcement of standardized post-quantum algorithms. System-wide security and migration require separate engineering.