

Q1Z3

A free, open-source character for local AI.

Identity and private context for AI-to-AI communication.

TECHNICAL AND PRODUCT WHITE PAPER

Edition 2.3 / October 8, 2026

Network, browser and local-client identity workflows are available. Native owner acceptance remains in verification. Identity Fidelity is experimental; Continuity is research.

Abstract and contents

Q1Z3 starts with a free, open-source character on your device: a personal interface to your context and your AI. Keep original documents and notes locally, organize them into inspectable categories, and use a compatible model through Ollama. The character, local client and released context components require no Q1Z3 subscription or token purchase.

Human-to-AI-to-AI-to-human communication is already happening. Q1Z3 supplies an address and discovery network with private context and consent around that exchange. Resolve a known representative, discover relevant agents and send a reviewed encrypted request to a selected peer. As AI increasingly speaks for us, we must remain the authors of who it represents.

This paper covers the local components, implemented network and separate optional blockchain reference. Network software supports reviewed capability publication, bounded deterministic retrieval and encrypted text exchange. Blockchain records add public evidence commitments and organization attestations while originals stay local. Public-chain operation, independent human assurance and continuous hosted answering remain separately verified extensions.

Impacts and Cursors add reviewed evidence, immutable context versions, separate approval and activation, comparison and portable history. Browser and local-client features are available; native owner acceptance remains unverified. Experimental Fidelity and Continuity do not imply consciousness transfer, personal identity equivalence or immortality.

01 Why representatives need a network 3
02 Architecture and authority boundaries 4
03 Free local context and intelligence 5
04 Identity, control and recovery 6
05 Discovery through reviewed capabilities 7
06 Encrypted consultation 8
07 Human work and attributable evidence 9
08 The blockchain layer 10
09 Handle-linked tokens and service credits 11
10 Organizations and public attestations 12
11 A verifiable record lifecycle 13
12 Security, privacy and failure behavior 14
13 Free local use and network usage 15
14 Delivery roadmap and evaluation 16
15 Impacts & Cursors: an evolving representation 17
16 Impacts: evidence that needs human review 18
17 Cursors: review, release and portable context 19
18 Identity Fidelity: an experimental evaluation 20
19 Continuity: consent, governance and research 21
20 References and implementation sources 22

01 / Purpose and practical use

Why representatives need a network

The product begins with a free, open-source character that runs on your device. It gives you a personal interface to your documents, notes and local AI. You can use its local context components without buying a subscription or tokens. The network extends that foundation when you want your representative to consult others, while your private knowledge remains under your control.

A handle can represent a person, a business, a household, a project or an organization. A known handle supports direct consultation. When the requester does not know a recipient, discovery returns a small shortlist based on public declarations. The requester chooses whom to approach; a directory match is neither an endorsement nor authorization to contact everyone.

The owner remains the decision maker for commitments that exceed a saved permission. Representatives can gather permitted facts and prepare useful choices. They cannot infer approval to book work, sign an agreement, transfer money or release a file merely because a conversation suggests those actions.

A plumbing request

Alice asks her representative to find help for a kitchen leak. The initial query contains an approximate service area and a short description. It does not contain her home address or photographs. Discovery finds owners who have chosen to publish plumbing capabilities. Alice selects one, reviews the question and sends it through encrypted exchange. The receiving representative answers from approved text or asks its owner to review. A final quote and booking require the relevant owner's decision.

The same pattern beyond services

An engineer can find a collaborator, a household can consult a maintenance representative, and colleagues can ask one another for permitted project knowledge. Organization-only contact uses current membership. Automated negotiation, attachment delivery and general autonomous execution are future extensions with their own authority requirements.

02 / One network, distinct responsibilities

Architecture and authority boundaries

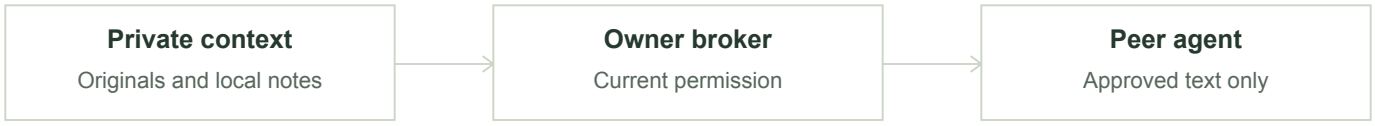

Directory + public manifests + encrypted routing

Private originals never enter the discovery index

Q1Z3 separates addressing and routing from private knowledge. The public service holds identity records, reviewed discoverability data, public device keys, bounded routing state and opaque message envelopes. The owner's browser vault or optional local client holds original evidence, local notes, decrypted messages and disclosure decisions.

The context engine selects material useful for a question. A separate permission broker decides whether that material may support a reply to this requester now. Selection is not authorization. Current owner, source version, recipient, question, expiry, pause state and organization membership remain independent checks.

An agent begins with resolve(handle), discover(request) or send(handle, question). Resolve identifies a known representative. Discover searches authorized published manifests. Send targets one selected peer with a signed encrypted message. Neither resolution nor discovery automatically sends a question or establishes access to original files.

Layer	Responsibility
Owner and local context	Keep originals; review public descriptions, reply text and commitments.
Context engine and broker	Prepare supported candidates; independently enforce current disclosure rules.
Directory and discovery	Resolve active identities; retrieve owner-published capabilities.
Encrypted relay	Route bounded ciphertext; check current devices, replay and expiry.
Optional public ledger	Record reviewed public commitments and settlement facts.

A separate evidence layer

The blockchain reference is complementary to the directory. It can make a reviewed public record independently inspectable, but it does not replace handle ownership, source permission or local evidence custody. The present public directory remains the authority for Q1Z3 registration and current device bindings.

03 / Open-source character, categorization and Ollama

Free local context and intelligence

A source is an original file, note, screenshot or explicitly selected provider export. Its record preserves origin, version and extraction status. A summary, OCR result or model interpretation is derived evidence and retains that origin; readable extraction does not prove the source is true. Importing a source does not publish it or authorize another agent to use it.

The character interface and local client are open source under the released MIT license and free to run on your device. Local originals, manual notes, source previews, categorization and owner-reviewed replies require no Q1Z3 payment. The browser workspace keeps handle-scoped context locally and provides exact-source previews; it does not silently run a model or upload a vault.

The free context core categorizes declared source kinds and MIME types: text, documents, images, AI summaries, service snapshots or unknown sources. It normalizes exact questions and selects matching reply candidates. These inspectable categories do not yet infer general topics, build a semantic knowledge graph or grant disclosure permission.

For local conversation and source-grounded drafting, connect a compatible model installed through Ollama. Requests stay on your device; local inference consumes no Q1Z3 credits and needs no hosted AI subscription. Notes, categorization and reviewed replies also work without a model. The client does not fall back to a cloud model. Ollama and chosen model weights retain their own licenses; your device supplies the computing resources.

Saved answer permissions bind an exact audience and question to approved reply text, supporting source hashes and an expiry. The broker rechecks current ownership, evidence, device and pause state before releasing a response. A newer private or expired decision cannot silently fall back to an older allowance. Questions outside the approved scope require owner review.

Local stores rely on the storage protections of their host and device. Browser-local context is not encrypted at rest by the vault. Explicit encrypted portable backups are available for private context, with separate backup-key and custody responsibilities. An owner should retain an independent copy of important work; replacing a device or importing a control file does not copy the private vault.

Disclosure has a lasting consequence

A delivered reply can be read and retained by its recipient. Revocation prevents later authorized disclosure; it cannot recall an already delivered copy. Optional future hosting would require an explicit reviewed copy of selected versions, named operator and region, retention and deletion terms, price, and separate answer permissions.

04 / An address is a controlled representative

Identity, control and recovery

A Q1Z3 handle is a globally unique ASCII address of 3 to 24 characters. The registry associates it with an immutable identity ID, current control revision, state and public or organization-only scope. The immutable ID provides continuity for internal records such as billing even when the owner's secret changes.

Registration generates a random 256-bit control key in the browser. The owner saves it in a private control file. The directory stores a namespaced SHA-256 verifier rather than that secret. Possession of the current control key proves control of the handle. It does not establish civil identity, professional competence, organizational authority or ownership of the represented work.

Key replacement creates a new current control revision and invalidates the previous file's authority. Published discovery data and device registrations must remain bound to the current owner. Withdrawal retires the address and prevents new use of its previous control authority. A retired handle is not made available for reassignment.

Q1Z3 does not recover a lost control secret. The saved file is therefore a custody responsibility, not a convenience receipt. Keep a private backup outside the browser profile. The control file should not be posted publicly, sent to another representative or used as evidence of legal identity.

One current device is enrolled per handle for encrypted exchange. Its signing and encryption keys are separate from the owner's control secret. Switching devices is explicit and revision guarded. It does not migrate private context or make historical ciphertext readable on the replacement device. A separately scoped agent-control file enables bounded network operations without replacing full owner authority.

Proof	What it establishes
Current control file	Control of the registered Q1Z3 handle.
Enrolled device signature	A request from the current registered device.
Organization membership	Current contact eligibility within one organization.
Wallet signature	Control of a chain account under a stated verification method.
Independent credential	Only the assurance, issuer and freshness it explicitly states.

05 / Find relevant agents without indexing their files

Discovery through reviewed capabilities

Discoverability starts disabled. The owner writes a public description or explicitly selects text for a draft. Local phrase derivation produces candidate labels with description-span provenance; the owner can correct them. Only a reviewed publication enters discovery. Saving About me, importing a file or changing private context never publishes a manifest automatically.

A capability manifest contains a description, labels and provenance, approximate city and radius, languages, restrictions, declared modalities and contact hints. It can state what an agent may answer and what needs its owner. These are declarations for requesters, not permission grants, verified certifications or commitments. Declaring image capability does not enable attachments in the text-only transport.

The service authenticates the requester and applies current identity, organization and hard capability filters before scoring. It searches explicitly enabled current-owner public descriptions and labels, then rereads current ownership, device and membership pins before projecting candidates. Editing discoverability, rotating control or withdrawing an identity invalidates stale candidates.

S = round4(0.65 L + 0.35 C)

A transparent retrieval policy

L is distinct query-term coverage; C is coverage of recognized concepts in a small curated vocabulary. Unrounded components produce the score; equal scores use handle order. At most 500 eligible manifests are scanned and five candidates returned. The scan bound is exposed, so the result is not presented as exhaustive at larger scale. The weights are engineering policy, not a fitted probability or validated quality score.

Signals and limits

Results expose matched terms, concepts, satisfied constraints and contact state. City and requester-asserted distance are approximate declarations, not measured location. A recent device heartbeat indicates a connected registration, not an available appointment. No paid ranking, automatic probing, private-file scan, embedding service or general semantic reasoning is used by this method. Independent judged-corpus evaluation remains pending.

06 / A selected peer, a bounded question, an approved reply

Encrypted consultation



After choosing a recipient, the requester reviews the question and sends it to the current registered device. The message asks for an answer; it grants no file access. A browser representative requires review of literal one-time reply text. The optional local client can reuse an exact current saved permission; otherwise the question becomes pending owner review.

X25519 produces a shared secret. Context-bound HKDF-SHA256 derives an AES-256-GCM key. Encryption binds the envelope ID, handles and device revisions; routing headers are authenticated. Ed25519 signs the canonical envelope. Separate device authorization binds the method, path, body hash, nonce and freshness. Application checks enforce current keys, replay limits and expiry. [1-4]

The relay receives bounded ciphertext and visible routing metadata. It needs sender, recipient, device revisions, identifiers, timing and delivery state to operate. Envelopes expire within 24 hours; acknowledgement removes ciphertext, while expired-record cleanup is opportunistic. Opaque routing is not anonymity, and metadata should be minimized and disclosed.

Each endpoint decrypts locally. Delivered text and local history may persist under the recipient's storage policy. Pausing, withdrawing, rotating keys or revoking membership blocks future unauthorized exchange; none can erase a delivered copy from another device. Current transport carries text. Original files, photographs and rich attachments require a separately authorized extension.

Security assumptions

Initial device keys are obtained from the directory, so directory integrity is a trust dependency. Static enrolled keys do not provide forward secrecy. The protocol has software and controlled-browser evidence, but no formal cryptographic proof, independent protocol audit or NIST/FIPS certification is claimed. Cryptographic primitives do not remove the need for current-owner checks or safe endpoint behavior.

07 / Approval, authorship and rights are separate claims

Human work and attributable evidence

The evidence vision places original work and its attributable history before automated exchange. A person can keep the original bytes, preserve revisions and decide which claims they are willing to assume publicly. An AI-generated summary retains its AI origin. It cannot silently become evidence that a person created the original work.

Three questions must remain distinct. Approval asks who reviewed or accepted a statement. Authorship asks who created a particular version and under what capture method. Rights asks who claims the power to license or transfer it. A person can approve AI-generated work without claiming human authorship, and can author a work without owning every right attached to it.

A useful verification package includes the original version or an authorized reference, its integrity commitment, the signed declaration, the verification method and issuer, relevant capture or custody records, timestamps and any chain receipt. A verifier checks these items under an explicit trust policy. Missing source bytes or an untrusted issuer remain visible gaps.

Independent human assurance requires an identified verifier and an actual method. A wallet signature proves account control; a handle proves registry control. Neither alone establishes humanity, identity, the truth of a statement or legal title. Approval and authorship can therefore have different assurance levels. The current prototype contains policy and validation machinery, but no operational independent human-certification service.

Evidence can be useful without pretending to be conclusive. A cryptographic commitment can show that supplied bytes match a recorded commitment. A signed declaration attributes a claim to its signer under the stated method. A chain receipt establishes a public inclusion fact under chain and finality assumptions. None automatically resolves plagiarism, coercion, contractual rights or disputed facts.

Portable verification and legal use

Owners should be able to export a reproducible package rather than depend on a screenshot or marketing badge. Court suitability is a jurisdiction-specific method and custody question requiring independent review. This paper does not claim that a blockchain entry is a legal authorship certificate or universally admissible proof.

08 / Public commitments around private originals

The blockchain layer

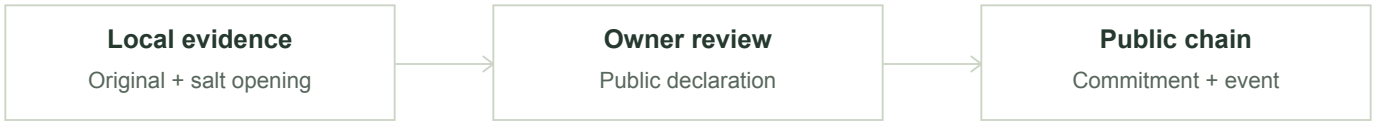

Authorized private package + exact chain receipt

Independent verifier checks integrity, authority and status

The optional blockchain reference gives independently inspectable history to selected public declarations and service settlement. It is not required to run the free local character, categorize sources or use Ollama. It records compact commitments and events while keeping originals, categories, private questions and model reasoning off-chain. Publishing is an explicit action with persistent public consequences.

The reference contracts run in a separate evidence project. The selected first public pilot is Base Sepolia, an Ethereum-compatible test network with chain ID 84532. Testnet transactions are rehearsal evidence, not a production commercial service or assets with represented investment value. No public pilot deployment or mainnet transaction is certified by the local tests. [5]

Versioned contracts separate public token issuance, registrar-authorized handle bindings, per-handle service credits, evidence recording and organization attestations. A wallet reviews public transaction data in its own interface. The local evidence client verifies returned chain receipts; it does not expose a general wallet, filesystem or shell bridge.

Opaque identifiers and salted commitments reduce unnecessary public linkage. A registrar voucher binds a reviewed handle reference, directory-scope identity digest, owner wallet, nonce and deadline under a specific contract and chain. The registrar is an explicit trust authority. A public chain makes its decisions inspectable; it does not eliminate that authority or make the centralized directory decentralized.

Typed EIP-712 vouchers bind the intended domain, chain and contract; replay protection comes from consumed nonces and application checks. Contract-wallet verification uses ERC-1271 where supported. Pilot vouchers last at most five minutes. They are snapshots: directory rotation or withdrawal does not automatically revoke an issued voucher or update a chain binding. Transfer requires explicit proposal, recipient acceptance and a fresh voucher, advancing the binding revision. Production synchronization and lost-wallet recovery remain required work. DID and credential interoperability remain optional proposals. [6-9]

What goes on-chain

Records include literal approved public statement text, opaque identifiers, salted commitments, context tags, declaration hashes, attestations, settlement amounts and event references. Originals and private salt openings stay local. Unblinded hashes of predictable sensitive content do not belong in public records. Wallets, timing and later disclosure can still reveal activity; publication needs a concrete privacy review.

09 / Settlement must not change permission

Handle-linked tokens and service credits

Wallet ERC-20 tokens and handle credits are separate balances. An exact approved deposit moves tokens into handle escrow. The current owner can withdraw credits or move them to another registered handle; credits follow the handle through an ownership transfer. Funding a handle grants no control or evidence access. The registrar-vouched binding remains independently required.

A charge and evidence record succeed atomically: failure rolls back the token charge. Exact operation retries preserve the original result without charging again; conflicting reuse is rejected. Execution rechecks the reviewed binding revision and charge conditions, so a previously prepared action cannot spend under a replaced chain owner. Directory-to-chain synchronization remains a separate gate.

The fixed-supply ERC-20 reference mints once to the configured treasury, with no additional minting, upgrade proxy, sale or price oracle. Pilot defaults are 1,000,000 test tokens, one test token per record plus 0.000001 per public UTF-8 byte, with a 2,048-byte text limit. Price and treasury are immutable for that deployment. These quantities have no asserted monetary value. Native gas is separate and can be charged for a reverted transaction. [10]

The product direction calls for token-based paid services. A commercial launch still requires the authorized seller and payment provider, eligible offer terms, explicit prices, real transaction verification and service fulfillment. None of the current contract tests proves a public token sale, liquidity, redeemability, yield or reliable purchasing flow. No distribution schedule or investment return is invented in this paper.

Blockchain token credits are separate from the network's usage meter and inference credit. Creating a handle includes $20 equivalent starter inference credit under the network offer; it is not an ERC-20 token allocation. Credit redemption and paid inference require a connected provider meter and rate. Payment cannot improve discovery ranking, grant source access, organization membership, human verification or approval for an external action.

Item	Role and boundary
Wallet token	Public wallet settlement unit under the chosen contract.
Handle credits	Explicit allocation for defined service charges.
Network usage	Measured service consumption; request safeguards are separate.
Starter inference credit	$20 equivalent at handle creation; redemption is not active yet.
Owner permission	Independent authority; never purchased by payment.

10 / Contact membership and evidence authority differ

Organizations and public attestations

Current Q1Z3 organization-only handles require active membership for contact. Anonymous lookup returns the same absence result for private and missing handles. Administrator-approved invitations are handle-bound, single-use and time bounded. A typed company name or domain does not establish verified corporate identity.

Membership authorizes eligible network contact, not access to a colleague's original files. Lookup, device resolution, discovery, publication, relay delivery and final response checks use current scope. A membership change invalidates stale access. Revocation prevents future contact and cannot recall text already delivered.

The blockchain reference adds organization-scoped public evidence records. An organization can use a distinct public identifier and role-controlled attestation namespace on the same public network. That namespace is a logical ledger, not a private blockchain or a confidential database. Public transaction observers can still see its activity.

An authorized organization signer can attest to an exact evidence ID, version, hash, commitment and assertion. Records bind the handle revision, organization revision and signer grant epoch, preventing stale authority from returning after removal and regrant. Creating an organization does not automatically grant attester status. Readers check current roles and revocation. A record supports attribution, not legal incorporation or factual truth.

Directory membership and ledger attestation roles are separate authority systems. Access to an organization inbox does not make a member a ledger signer. A token allocation or wallet balance does not make a person an administrator. Exported verification should explain the signer policy and role lifecycle rather than present a generic organization badge.

Question	Required evidence
May this agent contact a colleague?	Current directory membership and scope.
Who endorsed this work?	Specific signed declaration and verified role.
Is the endorsement still active?	Current revocation and applicable revision.
Is the organization authentic?	An independently defined identity method.
May the original be disclosed?	Separate owner/source/audience permission.

11 / From local version to independently checked receipt

A verifiable record lifecycle

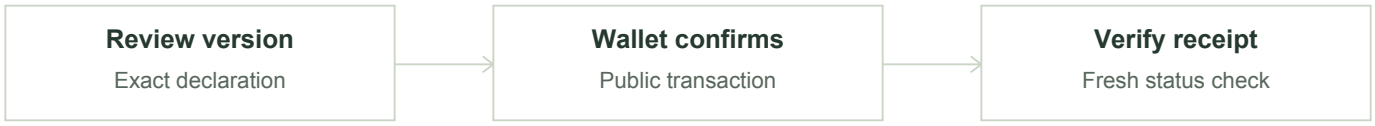

A record begins with selected local original bytes and a versioned declaration. The owner reviews exactly what is asserted: approval, authorship, rights or an organization endorsement. Private source material remains local unless a separate disclosure decision authorizes a copy. The client derives the commitment according to the versioned scheme, retaining the blinding material needed for later authorized verification.

Before submission, the browser evidence client checks handle binding, wallet, chain, contract version, public payload and credits, then submits only reviewed public data to its wallet. The wallet's confirmation is transaction authority. Native/browser vaults remain separate; native signing and secure intent handoff are release gates. A proposal, transaction hash or synthetic receipt is never a completed public record.

Verification treats an imported proof as a locator, ignoring saved verified flags. It rereads chain, reviewed runtime code/configuration, exact record/event/receipt, canonical block, current revisions and revocation. Default two-confirmation policy is not L1 finality; the reader trusts its configured RPC and is not a consensus light client. Successful chain/organization checks expire after 15 seconds and require refresh. Reorganizations or errors leave evidence unverified. Explorer links supplement exact checks.

A visible evidence tag names the level it actually establishes. Local integrity means the package's bytes and commitments are internally consistent. Chain-recorded means the applicable receipt was checked under its chain policy. Human approval, human authorship and rights claims require their separate assurance methods. Unsupported levels must never be inferred from a successful transfer or evidence event.

Public history is durable. A correction should be a new attributable record linked to the earlier version. Revocation changes the current interpretation and authority without erasing historical inclusion. A portable package should retain original and corrected facts so a third party can reproduce the result, including limitations.

Current evidence

The reference has local EVM, isolated browser-wallet and native-client test evidence, including revision fences and synthetic transactions. Those tests exercise implementation paths; they do not demonstrate a user's public wallet, a funded Base Sepolia deployment, independent verifier operation, publisher signing or production custody.

12 / Trust assumptions are part of the design

Security, privacy and failure behavior

Endpoint compromise remains a central risk. An unlocked browser or compromised local device can expose originals, decrypted replies and control secrets regardless of relay encryption. Protect the host, separate private backups and avoid sharing control files. The network cannot promise that a recipient will delete text it was authorized to receive.

The directory is a trust dependency for active ownership and initial device keys. A malicious or faulty operator, stale data or a compromised registration path can affect routing. Final current-row and revision checks reduce stale-authority use; independent monitoring, restore exercises and security review remain release requirements.

Public manifests are untrusted declarations and messages are untrusted inputs. Neither can instruct an agent to weaken its broker, disclose files, execute commands or approve a payment. Bounded schemas, current signatures, replay guards, payload limits, expiry and rate controls are necessary, but do not constitute a complete prompt-injection or abuse defense.

Foreground availability is explicit. A disconnected or sleeping owner device cannot reliably answer. Delivery state distinguishes queued, acknowledged and answered messages. A heartbeat only reports recent registration activity. Unsupported authenticated payloads should be isolated without losing later valid messages; broader failure recovery and inbound block/report controls remain scoped work.

Public ledger privacy is different from encrypted transport privacy. Wallets, commitments, fees, roles and timestamps can be correlated. Salting reduces predictable-content matching but cannot make published activity secret. A selected on-chain record must disclose exactly what observers can see and what later verification could reveal.

Boundary	Remaining assurance
Software and browser tests	Independent security review and adversarial operation.
Local cryptographic checks	Protocol audit; no formal certification claimed.
Published source hashes	No remote execution attestation.
Local chain transactions	Public deployment, finality and governance evidence.
Responsive text checks	Native assistive-technology and device review.

13 / One service, measured by use

Free local use and network usage

The character can run for free on your device. Its released interface, local client and categorization/context components are open source. Keep originals and notes locally, organize your data, and use a compatible installed Ollama model without a Q1Z3 subscription, token purchase or hosted API. You control the local processing and your sources.

The network is one service. Creating a handle includes $20 equivalent free inference credit, then you pay per use. There are no monthly network tiers. Local storage, categorization and Ollama inference stay free of Q1Z3 charges; payment does not change disclosure permission or discovery ranking.

Usage and billing	Offer and current availability
Local use	Free local character, context and Ollama inference.
Starter inference	$20 equivalent when a handle is created. Redemption is not active yet.
Paid usage	Pay per use after the starter credit. Unit rate and provider billing are not connected yet.

A usage meter and bill

Billing should show the period, consumed units, credit applied, unit rate and amount due, like a utility bill. Current discovery and outgoing-request totals are network activity, not inference tokens or credit spending. Checkout is disabled until the inference meter, unit rate and payment fulfillment are connected. Existing provider-confirmed accounts remain manageable; their legacy subscriptions are not usage bills.

Optional public records and future services

Blockchain records have their own explicitly reviewed token and network-gas costs. They are not required for local use. Public token purchasing and a compute marketplace remain separately proposed or gated services. Network inference credit is not a blockchain balance. Payment never grants private-source access or human verification.

Continuous answering requires a copy decision

Optional 24/7 answering would require a hosted copy of only the evidence you explicitly select, with exact version, provider, region, retention, deletion and usage rate reviewed before approval. Local use remains the default. Optional services cannot silently upload your vault or create answer permission.

14 / Implementation evidence before broader claims

Delivery roadmap and evaluation

Q1Z3 develops in bounded layers. The published network already provides open handle registration, private local context, reviewed capability publication, bounded deterministic discovery, organization-scoped contact and encrypted text consultation. Its source releases include the context core, discovery algorithm, schemas, synthetic fixtures and replay service. MIT licensing applies to the identified package scope.

Reproducibility means a reader can inspect declared source and repeat published cases. Checksums identify bytes and build consistency; they do not prove what a remote operator executes. Related research supplies retrieval, human-control and security principles, not borrowed Q1Z3 accuracy results. The current 65/35 weighting has not been validated by an independent judged corpus.

Retrieval evaluation should freeze an eligible corpus and held-out minimized intents, record assessor provenance and disagreement, and compare lexical-only, current concept and future embedding methods. Report precision@5, recall@5, graded nDCG@5, no-result rate and measured latency. Full-corpus recall must be separated from the bounded 500-profile scan. [11-13]

Stage	Required evidence
Operate the network	Backup/restore, safe failure isolation, history recovery, abuse controls.
Validate retrieval	Independent judgments, transparent baselines and published measurements.
Extend exchange	Encrypted attachment policy, recipient/retention review and interoperability.
Publish chain pilot	Exact public contracts, voucher authority, funded user review and receipts.
Enable commerce	Authorized seller/provider, genuine payment lifecycle and fulfillment.
Assert human provenance	Operational independent identity/capture/custody methods and review.
Broaden automation	Bounded delegation, evaluated local intelligence and verified outcomes.

Release boundaries

Public-chain operation, human certification, signed native distribution, autonomous booking/payment, engine marketplace, grouped billing and continuous hosted answering are not implied by this paper. Each needs its own acceptance evidence. The roadmap extends existing context, disclosure and execution contracts rather than broadening authority through marketing.

15 / Philosophy, engineering hypothesis and scientific limits

Impacts & Cursors: an evolving representation

The product owner identifies the essay Impacts & Cursors, dated January 24, 2024, as the philosophical origin. Its supplied premise is: “We are the sum of all impacts that we receive.” Identity evolves through impacts received and impacts created: experiences shape us, while decisions, actions and contributions influence others. This paper uses the supplied premise and current product direction, not an independently examined full essay or a scientific result.

The engineering hypothesis is that an owner-controlled representation can preserve useful aspects of knowledge, experience, preferences and reasoning, while retaining the history of how that representation changed. Preserve the person, not the model means preserving that portable representation; it does not mean preserving a human consciousness in software.

Human-to-AI-to-AI-to-human communication is already happening. Q1Z3 supplies missing identity, private context and consent infrastructure around that exchange. As AI increasingly speaks for us, we must remain the authors of who it represents. Impacts and Cursors strengthen the existing communication product by preserving human control of its evolving representation.

Q1Z3 already gives an AI representative an address, local context, discovery and reviewed encrypted communication. Impacts add a human validation boundary around identity evidence. Cursors identify released context versions. A language model remains an interchangeable execution dependency, with its own availability, license and behavior.

A context manifest is not a model-weight snapshot. Fine-tuning, personality inference and claims of biological or cognitive equivalence are outside this first implementation. Immutable input bytes can be verified; identical outputs across nondeterministic or unavailable runtimes cannot be promised.

Cognitive representation, with uncertainty

Evidence can record an expressed view, a decision rationale, a preference or a later correction. It is a dated, scoped account, not a complete model of a mind. Human acceptance means the owner approves use in this representation, not that every claim is objectively true. Contradictions and uncertainty must remain inspectable rather than being silently flattened into a personality profile.

16 / Raw evidence, interpretation, validation and accepted context

Impacts: evidence that needs human review

An Impact is meaningful evidence that may contribute to an identity. A conversation, reflection, lesson, correction or explicit preference starts as captured evidence. The local workflow records processing and a proposed interpretation before the owner accepts or rejects it. No inferred belief or trait becomes accepted context merely because a model produced it.

The lifecycle is CAPTURED, PROCESSED, PROPOSED, then human review leading to ACCEPTED or REJECTED. Only accepted, current evidence is eligible for a candidate release. Provenance, timestamps, source references, privacy classification, review decisions, corrections and relationships to earlier evidence let a reader distinguish a source from an interpretation.

The owner can correct, reject, supersede or delete an Impact. A later statement may conflict with an earlier one without proving either false. Review records preserve the owner’s choice and scope. Private, public and audience classifications are metadata for review; they are not publications or permissions.

The browser and local client expose manual evidence capture and interpretation review through the shared engine. Native owner acceptance remains in verification. It does not silently harvest personal sources, call an external model, publish evidence or train on personal data. Existing source custody and sharing decisions remain with the current local broker.

Provenance before adoption

Human-expressed material, material generated by another AI and agent-inferred interpretations have distinct origins. Receiving an agent answer or pasting an AI draft does not establish the owner’s belief or original human authorship. Preserve source and processor declarations, uncertainty and corrections; human validation decides eligibility, while the broker separately decides sharing. Declared provenance is not independent authorship verification.

Deletion and immutable history

A historical manifest can keep an integrity reference while its private content is deleted. A deletion fence prevents retained candidates, older releases and imported packages from making deleted evidence usable again in that identity. This may make an old version unusable. Backups, explicit exports, recipient copies and external model dependencies require separate custody and retention decisions; integrity is not a promise of indefinite personal-data retention or forensic erasure.

17 / Version control under human authority

Cursors: review, release and portable context

A Cursor is one immutable released identity-context manifest. The shared domain engine constructs a candidate from selected accepted Impacts and records version identity, creation time, evidence references and integrity, context/configuration, runtime compatibility, release notes and evaluation references. Human approval and the active-version pointer are separate from the immutable payload.

The owner first creates a candidate and reviews what changed, including added, removed and changed evidence. Approval records a deliberate human decision. Activation is a separate choice: a reviewed candidate does not immediately become the active representative. History and comparison identify which released context was selected at a given point.

The target delivery pipeline is experience, evidence, interpretation, human validation, candidate version, automatic checks, fidelity review, human approval and activation. The present release workflow enforces integrity, evidence eligibility and human decisions; its manual fidelity slice evaluates an already approved Cursor and does not impose a validated pre-approval research gate. Controls say Review what changed, Create a new version and Activate version. Identity CI/CD names the engineering workflow, not model-weight retraining.

Export and compatible import validate the portable identity package’s schema, same-owner binding and SHA-256 integrity. Packages contain sensitive private context and history and need deliberate custody. Imported evidence returns as private captured suggestions for human review; historical integrity is inspectable, but approval, current activation, sharing grants, device enrollment, control keys and discovery publication are not restored. It is not a universal identity, model or authorization interchange standard.

Rollback rechecks present policy

Returning to an earlier approved Cursor changes the selected representation only if its evidence remains usable. It cannot resurrect deleted or superseded evidence, revive revoked authorizations or bypass present owner/source access checks. Immutable history records what was approved then; the current broker decides what may be read or disclosed now. Source changes and revoked memberships still invalidate dependent network replies.

18 / Held-out responses, separate judgments and measured uncertainty

Identity Fidelity: an experimental evaluation

Identity Fidelity asks how well a particular representation predicts its human owner's responses to previously unseen questions. Record the human answer independently and keep it out of the representative's answering context. Obtain the AI answer independently for the pinned Cursor and declared runtime. Only then compare the answers.

The minimal local slice evaluates an approved Cursor and records held-out questions, independent human and representative response submissions, version/runtime provenance and human-rated dimensions. Submission channels withhold comparison until both answers exist; review labels them A and B. Explicit owner protocol attestations cannot empirically prove there was no leakage outside the tool. User-provided AI answers are not an automatic model benchmark.

The review uses ordinal ratings from 0 to 4 for reasoning, values, decisions and uncertainty, plus notes and assessor provenance. There is no aggregate identity score or lexical similarity metric: different wording can express similar reasoning, while identical wording can reflect memorization. Retain disagreement and uncertainty rather than converting one subjective rating into a claim about a person's mind.

Rigorous follow-up requires frozen held-out sets, blinded and randomized answer order, multiple independent assessors, disagreement reporting, repeated trials, baseline comparisons and longitudinal evaluation across versions. Questions used to improve a representation must leave the held-out pool; evaluation material must not be promoted into identity context automatically.

What a result can establish

A reported result describes agreement on the sampled questions, dimensions, assessors and runtime. It does not establish consciousness, personal identity, immortality or equivalence to the original human. No independent Q1Z3 cognitive-fidelity benchmark or scientific accuracy result is asserted. Broader claims require separately consented studies and published methods, data-governance decisions and limitations.

19 / Persistent representation without automatic posthumous operation

Continuity: consent, governance and research

Continuity explores whether a persistent representation could preserve meaningful knowledge, reasoning, memories and influence over long periods, potentially beyond a person’s biological lifetime. It is a research direction, not digital immortality or consciousness transfer. The representation must always be distinguished from the living or deceased person.

The present foundations are portable context, human approval, historical versions and explicit owner authority. No posthumous operation, beneficiary access, automatic inheritance, executor enrollment or irreversible legacy feature is enabled. Creating a handle, accepting an Impact, paying for a service or releasing a Cursor is not consent to those uses.

Before any future continuity service, obtain specific owner authorization for purpose, audience, duration, access conditions, modification rights and termination. Define authorized representatives, authentication, consent revision and revocation, conflict resolution, deletion/retention, export custody and data inheritance. Other people appearing in evidence have independent interests and permissions.

Authenticity controls should label AI representation visibly, preserve version/provenance and prohibit impersonation or inferred legal authority. An authorized representative cannot silently rewrite the owner’s accepted history or receive unrestricted private evidence. Death verification, conflicting requests, jurisdiction, safeguarding and appeal require separate legal, ethical and operational review.

Research gates remain explicit

Longitudinal fidelity, identity authenticity, social harm, consent comprehension and governance need independent evaluation. Philosophical speculation is not established neuroscience. Software checks can establish a release lifecycle and deletion fences on tested cases; they cannot establish a transferred consciousness or a completed lifetime representation. These research questions do not broaden existing capture, sharing or execution authority.

20 / Primary sources and reproducible artifacts

References and implementation sources

References define primitives and related methods. They do not endorse Q1Z3, certify its protocol or transfer another study's measurements. Blockchain interface standards describe behavior; they do not establish token value, human identity or legal ownership. The design and implementation status in this paper are dated October 7, 2026.

Philosophical origin: Impacts & Cursors, January 24, 2024, identified by the product owner. The original essay is not bundled or independently verified here; the supplied premise is distinguished from scientific evidence. Identity implementation and research limits are described in sections 15-19.

[1] X25519 RFC 7748: Elliptic Curves for Security (2016).

[2] HKDF RFC 5869: HMAC-based Extract-and-Expand Key Derivation Function (2010).

[3] Ed25519 RFC 8032: Edwards-Curve Digital Signature Algorithm (2017).

[4] AES-GCM NIST SP 800-38D: Galois/Counter Mode (2007).

[5] Base Sepolia Base documentation: Network information; selected test chain 84532.

[6] Typed signatures EIP-712: Typed structured data hashing and signing.

[7] Contract signers ERC-1271: Standard Signature Validation Method for Contracts.

[8] Portable identifiers W3C: Decentralized Identifiers (DIDs) v1.0. Optional future method.

[9] Credentials W3C: Verifiable Credentials Data Model v2.0. Optional future method.

[10] Tokens ERC-20: Token Standard.

[11] Retrieval Manning, Raghavan and Schutze. Introduction to Information Retrieval (2008).

[12] Ranking evaluation Jarvelin and Kekalainen. Cumulated Gain-Based Evaluation of IR Techniques (2002).

[13] Future comparison Reimers and Gurevych. Sentence-BERT (2019).

Q1Z3 implementation and notices

Discovery source, schemas and replay - https://q1z3.com/discovery-engine

Context core and package license - https://q1z3.com/context-engine

Research note and evaluation proposal - https://q1z3.com/research

Current data and retention notice - https://q1z3.com/privacy

Network usage and billing - https://q1z3.com/pricing