

Understand control. Investigate change. Follow the evidence.

EXIT connects blockchain questions to source-linked answers, temporal context and explicit limits—building toward a research infrastructure for people, organizations and software.

Our first specialization is control and change analysis. Our broader direction connects authority, events and observed fund relationships across validated network coverage.

DOCUMENT STATUS · PUBLIC DRAFT 0.6

This document combines an existing engineering foundation with an approved development direction. Architecture diagrams distinguish delivered foundations from proposed extensions. Illustrative examples are not production benchmarks.

The public desk contains recorded cases. The same free API offers bounded Solana and Ethereum reads, research records, saved comparisons and source navigation. EXITP was created on 25 October 2026 at 12:14:57 UTC; free API access opened at 12:38:20 UTC. Paid token redemption and scheduled alerts are not live. The API Pass remains a private prototype requiring final-configuration chain rehearsal and security acceptance. Status correction, 27 October: initial delivery and Q1–Q4 belong to one continuing development program. Edition 0.5 is superseded historical material.

01 · Problem and mission

02 · How research becomes an answer

03 · Evidence and time

04 · Research graph and authority analysis

05 · Question planning and AI boundaries

06 · Monitoring and measured anomalies

07 · Delivery and verification

08 · Service rights and token integration

09 · First release and post-token quarters

10 · Advanced research and evaluation

11 · Security, privacy and transparency

12 · References and document scope

01 / PURPOSE

Public activity still requires accountable interpretation.

A transaction records an event. Understanding the event requires its authority, execution conditions, historical context and supporting sources. A current snapshot cannot, by itself, explain an entire history. An address does not, by itself, establish a person's identity or intent.

EXIT's mission is to make blockchain activity understandable, investigable and independently assessable. A result should help a reader determine what is supported, what changed and which observation would resolve an open question.


The central question

Who can do what, under which conditions, at what time—and what evidence supports that conclusion?

Security and operations teams can examine control paths. Researchers can reconstruct supported events. Developers can integrate structured answers. Individual users can inspect the same evidence through a guided interface.

From a question to an inspectable research result.

The engine evaluates supported questions against identified evidence. Planned research tools extend acquisition, navigation and investigation without replacing the evidence boundary.



Solid boundary: existing foundation, with bounded coverage. Dashed boundary: proposed development.

Figure 1. Conceptual architecture—not a map of private infrastructure or an assertion that every component is publicly available. Coverage and operating status must accompany any delivered release.

A source disagreement is retained rather than resolved by majority vote. Two endpoints may share an upstream dependency. Unknown history remains unknown even when the latest record is clear.

Preserve what happened and when it became known.

An event's effective time and the system's acquisition or knowledge time answer different questions. Later evidence must not silently appear in an earlier knowledge-time answer.

$r = (\text{subject}, \text{relation}, \text{object}, I_{\text{valid}}, I_{\text{known}}, \text{evidence}, \text{status})$

Conceptual relation representation. The exact implementation may use more than one record. This is an explanatory model, not a new published API schema.

$A(q, t_{\text{valid}}, t_{\text{known}}) = F_{\text{version}}(q, \text{eligible evidence}, \text{declared policy})$

Eligibility includes temporal visibility, source binding and coverage. A hash identifies the evaluated content; it does not establish the truth of all real-world assertions.



Figure 2. Synthetic example, not a measurement. Where the event time is not observed, EXIT must state an interval or an unknown start rather than invent an exact time.

Connect facts without inventing identity.

The expanded graph is planned to represent typed relationships among accounts, assets, contracts, roles and observed events. Every supported edge needs its source, time and interpretation boundary. A transfer, common funder or shared program is not automatically common ownership.




Figure 3. Synthetic example, not a measurement. Observed configuration is different from key custody, signing availability and the full set of execution conditions.

Threshold satisfied $\Leftrightarrow$ `|valid signatures from eligible members|  $\geq$  2`

Illustrative configured threshold only. This does not prove that two independent people exist, that their keys are available or that all program conditions permit execution.

Make the research plan inspectable before trusting the answer.

A proposed planner converts a user question into a bounded task: exact subject identity, network, time axes, required evidence and allowed cost. Missing identity or unsupported coverage must be resolved or reported before analysis proceeds.

```

plan = interpret(question, subject, network, time_scope)
if identity_is_ambiguous(plan):
    return clarification_required
if not covered(plan):
    return declared_research_scope_and_limits
records = acquire_or_reuse_permitted_evidence(plan, budget)
validated = check_sources_times_conflicts(records)
result = evaluate_supported_rules(plan, validated)
return result.with_sources_versions_and_unknowns()

```

AI may propose the plan and explain a supported result. It must not create a source observation, convert a hypothesis into a fact or initiate an unapproved financial action. The pseudocode describes a future product boundary, not an operational endpoint delivered by this document.

06 / MEASUREMENT

Investigate anomalies with visible error rates.

A planned Market Integrity module examines bounded classes of transaction, liquidity and coordination patterns. Research signals prioritize investigation; they do not automatically establish manipulation, responsibility or investment safety.

Precision

$$TP / (TP + FP)$$

Of the labelled alerts, how many matched the reference event?

Recall

$$TP / (TP + FN)$$

Of the labelled reference events, how many were detected?

Illustrative calculation: 18 true positives, 2 false positives and 3 missed labelled events produce 90% precision and approximately 85.7% recall. Unlabelled events are excluded and disclosed. These numbers demonstrate the calculation; they are not EXIT results.

Evaluation should freeze the event definition and reference set before the run, keep a held-out period, expose false positives and misses, and test source outages, changing market conditions and acquisition gaps. Historical replay does not guarantee future predictive performance.

One evidence foundation, three delivery paths.

Path	Practical use	Required boundary
Research interface	Ask a supported question, inspect the result and follow its sources.	Coverage, historical versus newly computed status and uncertainty remain visible.
API	Integrate structured answers into an application or organizational workflow.	Versioned contract, authentication, quotas, operating profile and support conditions.
Report / evidence package	Retain or exchange a reviewable investigation.	Only permitted content is shared; verification states its actual scope.

A reproducible output binds the question, evidence identities, evaluation version, time scope, result and limits. Recomputing a result checks those bindings; it does not create missing evidence or redistribution rights.

The current public research desk routes supported questions to recorded cases. It does not perform unrestricted live investigation or silently submit a research order.

A token right must correspond to a working service.

The selected design under evaluation is an optional API Pass: a supported token redemption buys quoted, bounded service credits. Normal service access remains available. Buying or holding a token alone does not establish an unlimited subscription or an unspecified financial entitlement.



Proposed production integration. Not deployed. Payment never changes the evidence conclusion.

Figure 4. Proposed production integration, not deployed. Service entitlement is a separate boundary from the read-only research engine; a private tested candidate is not a launched service.

For each eligible receipt q : number of credit grants(q) ≤ 1

Necessary invariant; insufficient by itself. An accepted finalized payment must also be recoverably credited under its quote. Network loss, restart and concurrency must not create duplicate or silently lost entitlements.

Token program, mint, distributions, treasury policy, pricing and venue conditions must be accepted on their actual configuration before launch. No supply allocation, token price or guaranteed fee income is invented here. Token creation, wallet approvals, signatures and the final launch action are performed with EXIT's creator wallet; no automated system can create or sign.

08A / CONFIGURATION & DISCLOSURE

A usable service first. An explicit launch configuration next.

EXIT is the project name. EXITP is the project token, created on 25 October 2026 through Jupiter Studio (Metora Dynamic Bonding Curve) on Solana. Official mint:

5vPz2s8NtkasmVRDz5ySz27V1E8UdSpzqc9xN678jups, read back from the chain as a standard SPL Token with 6 decimals, 1,000,000,000 tokens, mint authority null and freeze authority null.

Holding it does not currently redeem credits or activate a paid API Pass. The selected service design is an optional usage right; a token is not required to inspect the public research desk. Holding alone does not create an unlimited subscription, a profit share or an entitlement to operating revenue.

Proposed service principle

One credit corresponds to a scoped, completed research query with inspectable evidence, time coverage and remaining limitations. Technical errors, out-of-scope requests and unanswered primary questions do not consume credits. Known data gaps are disclosed before a paid query. A bare UNKNOWN is not a paid answer.

A private prototype enforces the stronger completed-answer boundary and records bounded follow-up intake without claiming that a research worker has executed it. Production adoption and alignment of all service descriptions remain release work; a local test is not public service availability.

Design parameter	Current position
Quote and credits	Accepted design: 15-minute quote, fixed integer credits, 12-month credit term, account-bound granted credits. Final prices and minimum redemption remain open.
Service continuity	Accepted design: no charge during technical failure; planned 90-day cessation notice and a receipt dispute process. Production operating evidence and final published customer terms are still required.
Payment independence	Payment does not change the answer, its confidence, correction or publication. A receipt cannot create a second credit grant when replayed.
Venue and token configuration	Venue: Jupiter Studio on the Meteora Dynamic Bonding Curve program (Solana). Read back from the chain after creation on 25 October 2026: SPL Token program, 6 decimals, 1,000,000,000 tokens, mint authority null, freeze authority null, no extensions; immutable metadata (name EXIT Protocol, symbol EXITP). Distribution, treasury and lock facts require their own dated disclosure.

EXPLICIT UNKNOWN / OPEN

Venue terms and issuer eligibility; prices, measured service capacity and minimum redemption; actual distribution and treasury/lock policy. No purchase, wallet action or launch is performed by publishing this document.

09 / DELIVERY & GROWTH

A finite first release. A continuing development program.

TOKEN CREATED · 25 OCTOBER 2026

Token creation occurred on 25 October 2026 at 12:14:57 UTC; free API access opened at 12:38:20 UTC. The first free usage right is live within its declared scope and quotas.

Token-to-credit redemption remains a separate future integration and is not activated by holding EXITP.

The quarters below do not add prerequisites to that first release.

T0 is the actual official token launch. Q1–Q4 are successive three-month planning windows after T0, not fixed calendar-year quarters. T0 is 25 October 2026 (the creation transaction, 12:14:57 UTC). Targets depend on completed dependencies, resources and accepted scope; the elapsed calendar does not establish delivery.



PLANNING WINDOWS · DELIVERY REQUIRES ITS OWN EVIDENCE

Actual 25 Oct sequence: EXITP 12:14:57 UTC; free API 12:38:20 UTC.

Post-token growth never becomes a new prerequisite for first release.

Figure 5. Relative post-token planning windows. These are future targets; elapsed time is not a completion receipt.

Q1

MONTHS 1–3

Research workflow and measurable service.

Saved research context, comparisons over existing coverage, clear source navigation and service cost/freshness baselines.

Method: reuse the query/export foundation; add typed plans and permission-bounded persistence. Evidence: pinned end-to-end tasks, negative controls and actual service measurements.

Q2

MONTHS 4–6

Temporal graph and selective network

expansion. Supported authority/event relations plus one bounded additional network or control surface selected by evidence, quality and cost.

Method: evidence-backed typed edges, chain-specific adapters and explicit bridge bindings. Evidence: golden graphs, wrong-join controls and finality/reorganization tests.

Q3

MONTHS 7–9

Measured monitoring and market investigation.

A bounded anomaly research module, replayable timelines and a report of detection quality.

Method: declared event classes, time-window features and held-out periods. Evidence: labelled precision/recall, misses, false alerts, outages and measured cost. Production real-time claims require their own operating proof.

Q4

MONTHS 10–12

Organizational workflows and advanced

verification. Permission-scoped workspaces, integration/report delivery and tested operational controls.

Conditional research: at most one bounded prototype in formal authority analysis or verifiable computation, only if resources and measured benefit support it. No promised production breakthrough.

Later expansion aims at broader validated cross-chain research, deeper scenario analysis and stronger machine-verifiable outputs. Each increment must improve a declared task; neither a larger budget nor an additional model proves user value.

09A / MEASURABLE EXITS

Progress is a tested result, not elapsed time.

Window	Required evidence for the selected increment	Stop or revise when
Q1	Preselected question-to-source tasks and missing/stale/conflicting-input controls pass; compare completion, errors and cost against the first-service baseline.	No improvement in the declared task, accuracy or cost; do not add an unrestricted planner to compensate.
Q2	Evidence-backed graph paths replay; zero wrong cross-chain joins in the declared hostile reference set; new scope meets its predeclared coverage/finality/cost profile.	The new adapter cannot establish identity, history or data permission within its budget.
Q3	Declare labels, evaluation window and thresholds before running the held-out reference period; report precision/recall with denominators, false positives, misses, unknown labels and latency/cost.	Quality is below that profile, labels cannot support the claim, or production latency/capture is unmeasured.
Q4	Cross-workspace isolation, revocation, restore and declared capacity pass; permitted evidence conclusions remain identical across payment modes.	An isolation/data-loss/security defect remains or the added workflow has no demonstrated benefit.

These are future increment acceptance rules. They add no prerequisite to the finite first release or the current token-readiness decision.

Use stronger methods where the task earns them.

Research direction	Potential benefit	Proof boundary
Formal authority reasoning	Analyze whether a modeled action is possible under declared contract and state conditions.	A proof applies to its model and assumptions, not unknown code or off-chain custody.
State and computation proofs	Verify supported data membership or a specified calculation without trusting a displayed answer.	Correct execution does not establish input completeness or economic causality.
Scenario analysis	Examine consequences of specific modeled authority or configuration changes.	A bounded technical scenario is not a forecast of market price or human behavior.

Adoption compares the proposed method against the existing baseline on the same tasks. Accuracy, unsupported claims, reproducibility, latency and cost are measured separately. A technique that does not improve the relevant outcome is not retained merely because it is fashionable.

Transparency without exposing private operations.

Public documents describe the evidence model, evaluation boundaries, coverage, releases and material limitations. They do not publish credentials, private investigations, internal host maps, withheld evidence or sensitive unfinished vulnerabilities.

Research and commerce data have distinct responsibilities. Access revocation, account isolation, durable accounting and recovery require adversarial testing. Wallet recovery phrases and private keys never enter the service.

The project may use a pseudonymous public identity while retaining honest operating and treasury disclosures. This does not remove transaction traceability or applicable platform requirements.

Software consumers can exercise technical tasks. Their results do not establish independent human usefulness, customer demand or legal eligibility. External participant recruitment is not a prerequisite for technical delivery.

Reference the method. State the implementation status.

This paper is authored for EXIT. External references inform its explanation and research directions; they do not imply endorsement, partnership or implementation of those systems.

1. Ethereum Whitepaper—example of a technical argument using state transitions, mechanisms and limitations. It is historical, not a current EXIT specification.
2. Solana: A new architecture for a high performance blockchain—example of system diagrams, explicit assumptions and analytical structure; its performance claims are not EXIT benchmarks.
3. W3C PROV-DM—provenance concepts for entities, activities, derivation and responsibility. EXIT compatibility is a possible development task, not a compliance claim.
4. EIP-1186—account/storage-proof proposal; network/provider support and trusted state roots must be evaluated.
5. Microsoft Research: Z3—symbolic reasoning tool considered for bounded future research.
6. RISC Zero zkVM—reference for verifiable execution, not a selected EXIT dependency.
7. GitHub public roadmap—reference for communicating feature stage and anticipated delivery separately.

Project source records distinguish existing local mechanisms, actual host results, public presentation and future development. This public draft does not claim that first-release or token-launch acceptance is complete. The illustrated architecture is explanatory; current access and release status are published separately.