TFA — Tela Filum Arachne

White paper, draft v0.2 · Sep 27, 2026 · @Thomas J Vaughn

Abstract

TFA is an independent network that observes supported blockchains, traces fund movements associated with flagged events, and preserves the evidence and its uncertainty so others can inspect it. Its first target is Solana, where suspicious token launches are common. TFA prioritizes transparent tracing after an event and analysis of links between launches: following funds hop by hop and connecting new launches to addresses involved in earlier flagged ones. TFA reports evidence and risk indicators, never accusations. It starts as a private observer and evidence ledger. It is designed to grow into its own blockchain that supports a native TFA coin, subject to a separate launch decision at the final stage. The name means web (Tela), thread (Filum) and spider (Arachne): the web of blockchains, the thread of the money, and the watcher.

The problem

On Solana, launching a new token is fast and cheap. That ease brings a steady flow of rug pulls: insiders take a large share of a new token, the public buys in, and the insiders sell or pull the liquidity, leaving buyers with a worthless coin.

Three gaps make this worse:

- **The money trail goes cold.** After a rug, proceeds split across fresh wallets, swap into other assets, and move to exchanges or across bridges. Ordinary users cannot follow it, and many tracing tools are paid products built for companies and law enforcement.
- **Scammers start fresh.** The same people can launch again under a new name with new wallets. Checks that look at one token at a time miss the link to earlier launches.
- **Findings are hard to verify.** Alerts and scores often come without the underlying evidence, so people must either trust them blindly or ignore them.

What TFA is

TFA is a watcher and an evidence keeper. It observes public blockchain data, detects measurable suspicious patterns, traces fund movements, and records what it saw with references anyone can check on the source chain.

What TFA does:

- Watches supported chains through one adapter per chain family, starting with Solana.
- Follows funds after a flagged event until they reach an exchange, a bridge, or go quiet.
- Links new token launches to addresses involved in earlier flagged launches.
- Keeps an append-only record of observations, assessments and corrections.
- Publishes findings through a read-only API and an explorer at tfachain.com.

What TFA does not do:

- It does not name people or declare anyone a criminal.
- It does not freeze, seize or recover funds. Exchanges, token issuers and law enforcement can act on its evidence; TFA cannot.
- It does not trade, hold user funds, or ask for private keys or seed phrases.
- It does not trace through privacy coins designed to hide transactions.
- It does not replace a legal investigation.

How TFA works

TFA turns raw blockchain data into published findings in a fixed order, and every finding keeps a link back to the transactions behind it.

Every TFA finding traces back to raw chain data

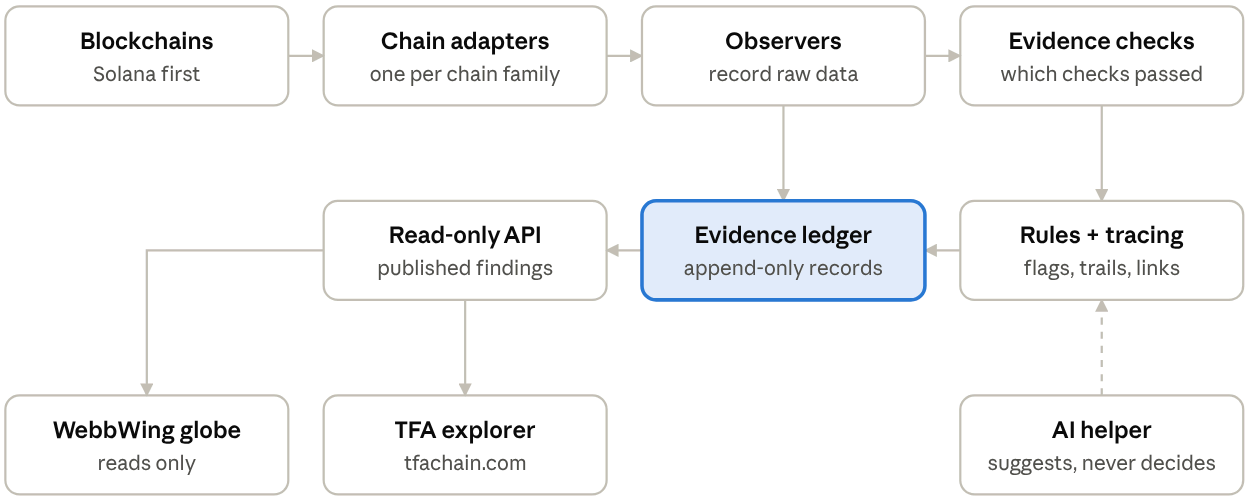

TFA architecture · from chain data to published findings

- **Chain adapters** translate each blockchain into one common format. Solana is first. Adapter code can be reused across compatible chain families, such as Ethereum-style chains, but each deployed chain still gets its own network identity, finality policy, list of supported programs and validation tests. Bitcoin follows. Each adapter documents what it supports and what it cannot see.
- **Observers** collect transactions and record which data source each record came from, so correlated sources are never counted as independent.
- **Evidence checks** record exactly which checks passed and which facts still rest on a provider's word.
- **Rules and tracing** raise flags, follow money trails and link launches. Every rule has a version, and every result names the rule that produced it.
- **The evidence ledger** stores records append-only and hash-linked. Changes are detectable against signed checkpoints kept separately and against independent copies; hash links alone cannot stop an operator from presenting a rewritten history. Corrections are new records, never silent edits.
- **The read-only API** serves published findings to TFA's explorer on tfachain.com and, later, to the WebbWing globe. Neither can change a finding.

If a data source goes down, TFA keeps its records, marks the gap in coverage, and never reports missing data as "no activity."

Coverage is published, not assumed. For every supported network, TFA publishes a coverage manifest: the network's identity, adapter version, supported transaction formats and program IDs, when observation started, available history, data sources, finality policy and known gaps. Coverage grows only after new programs pass tests.

Finality is tracked. Solana reports transactions as processed, confirmed or finalized. TFA records which state its source reported, treats processed data as provisional, and recomputes dependent findings if a source later reports something different.

Evidence model

TFA labels every statement with how it is known. These are separate properties of a record, not a ladder that ends in "verified fraud."

Level	What it means
Provider-reported	A data provider said the event happened.
Observer-attested	Observers using separately operated sources recorded the same event.
Checks passed	A named, versioned verification procedure succeeded.
Inferred	A rule or model concluded a relationship or risk from the facts above.

Each record carries its source transaction and slot, the data source used, a content hash of the raw evidence, the parser and rule versions, and its finality status. A hash proves the data has not changed since it was recorded; it does not prove the data was true to begin with.

Evidence is kept, not just its hash. A hash cannot rebuild a lost transaction record, so TFA defines what raw evidence it keeps, where, for how long, and how it is backed up and checked. When evidence becomes unavailable, TFA says so. Anyone reviewing a finding should be able to retrieve its evidence, verify the hash, rerun the named rule and get the same result. Private material, such as dispute correspondence, credentials and personal contact details, is never written to a public ledger.

Corrections. Mistakes are fixed by adding a correction record that names the record it replaces and why. The original stays visible. When a flag is withdrawn, every finding that depended on it is withdrawn or recomputed, so an error cannot keep spreading.

Public wording. Published reports state what was recorded and which rule fired, for example: "Recorded event X triggered rule Y, version Z, within the stated coverage." Reports about suspicious patterns carry the status **SUSPICIOUS ACTIVITY — NOT A DETERMINATION OF WRONGDOING**. TFA publishes no names of people. Each report shows a clear status: private draft, provisional, published, disputed, corrected or withdrawn. Anyone can dispute a report for free, without owning TFA coins; the process names who reviews it, deadlines, reasons for the outcome and an appeal. Removal of a valid finding is never for sale.

Following the money and linking repeat launchers

TFA prioritizes transparent tracing after an event and analysis of links between launches. How its coverage and performance compare with existing tools will be evaluated against documented alternatives.

Following the money (FLOW-001)

Starting from a flagged event, TFA builds a time-ordered graph of accounts, assets and transfers, with every edge linked to the transaction behind it.

What starts a trace. In the offline stage, a trace starts from clearly labeled test cases or a reviewed report with its source recorded. Before live monitoring, at least one fixed, protocol-specific starting rule must be approved: for example, a recorded liquidity withdrawal in a supported pool, with the amount, time window, the roles involved and innocent explanations. Its output is a measured event, not "rug pull confirmed." An outside accusation never silently becomes an established incident. Every rule has a written rule card: what it detects, supported programs, inputs, exact calculation, parameters, wording, innocent explanations, missing-data behavior, test cases, and who approved it.

- Each trace has stated limits (proposed starting defaults, not yet validated: 3 hops, 24 hours, 100 accounts), and every cut-off path says why it stopped.
- Exact transfers are shown as facts. Following a specific amount through a wallet that mixes funds is labeled as an assumption, and branches cannot count the same money twice.

- A swap counts only when a supported protocol's effects reconcile. Dollar values are dated estimates, never evidence.
- A trace stops at an unsupported bridge, missing history, its limits, or a labeled exchange address. Every address label has a source, a version and a way to dispute it. An exchange can identify its customer; TFA cannot and does not try. Cross-chain links count only when the specific bridge transfer is demonstrated, never because amounts look similar.
- "No observed movement within coverage" is reported instead of "funds did not move."

Linking repeat launchers (SOL-008)

TFA keeps launch roles separate: declared creator, signer, fee payer, mint authority, pool creator and funding address. It then reports two kinds of links:

1. **Exact role links:** the same address played a recorded role in two launches. TFA reports the role and the transaction, not "the same person."
2. **Funding-path links:** recorded transfers connect a new launch's role address to one in a previously flagged launch, within stated limits (proposed: 30 days back, up to 2 hops).

Shared exchanges, launch services, fee sponsors, routers and legitimate repeat teams are common innocent explanations, and every link lists them. Reusing an address proves only that the address was reused in that role, not that the same person controlled it or did anything wrong. Receiving funds never gives an address a permanent bad label, and a link to a launch whose flag was later withdrawn no longer counts.

Explaining a loss after it happens and warning before it happens are different abilities. TFA's first job is the first; it will not claim the second until it is measured.

Role of AI

AI helps TFA find and explain patterns, but it never decides what is true.

- AI may suggest unusual patterns, summarize a trace, and help prioritize alerts. Every AI statement must point to recorded evidence.
- AI output is stored separately from observations, with the model version and inputs recorded.
- A pattern AI suggests becomes a rule only after human review, with a written rule card, a version and false-positive tests.

- Agreement between several AI models is not treated as proof, because models share blind spots. TFA shows no AI "confidence percentages" unless a labeled evaluation supports them.
- AI runs isolated: it has no network access, no signing keys and no admin powers, so instructions hidden in token names or metadata cannot make it act.

Isolation limits what a compromised model can do, but it does not make its summaries accurate. So AI output is checked claim by claim against evidence and reviewed before anything is published. TFA's observation and rule processing must work fully with AI switched off. Connecting to an outside hosted AI model would change this design and needs separate approval.

The TFA network

TFA's long-term form is its own independent blockchain, run by many separate operators, whose ledger does not depend on any chain it watches.

Anyone will be able to download TFA node software and run it on their own computer or server. Planned node roles:

Role	Job
Watcher	Observes a supported chain and records what happens.
Checker	Re-checks other nodes' reports using a different data source.
Hunter	Flags suspicious activity first, with evidence.
Keeper	Stores copies of TFA's records so they cannot be lost.
Validator	Adds new blocks to the TFA ledger.

Independence is measured, not assumed. TFA tracks shared ownership, who controls signing keys, shared hosting and shared data sources; unknown independence stays unknown. Several nodes run by one person, or nodes that share one data provider, do not count as independent confirmation. One operator may perform several roles, but those roles are not separate votes.

Agreeing on the ledger is a different job from checking other chains. TFA's consensus decides the order and content of TFA's own records. It cannot make an incorrect Solana report true. The rules that update TFA's ledger must give the same result for the same inputs and never depend on AI output or live web responses. TFA will evaluate an established consensus implementation before inventing one; CometBFT is one candidate under study, not a decision.

Early multi-node testing will use a disclosed, permissioned set of validators, and TFA will call it permissioned. Open participation needs its own reviewed design for admission, incentives and resistance to one person posing as many; it does not follow from simply publishing node software. If a watched chain goes down, TFA pauses coverage of that chain and keeps its own history.

Until then, TFA runs as a private observer on the founder's servers, separate from any other service. A local hash-linked ledger is an audit log, not decentralized consensus, and TFA will describe it that way.

The TFA coin

TFA is designed to support a native coin, also called TFA (ticker TFA), created by its own network, subject to the Stage 6 launch decision. There is no planned sale and no promised price. All amounts and launch parameters are undecided.

How it would be earned. In TFA, "mining" means running a node and being rewarded for narrowly defined work that can be checked. Candidate categories, all to be tested first: assigned verification tasks, reproducible corrections of wrong records, and proven evidence-storage or archive services. Each task has a defined input, an authorized submitter, a deadline, an acceptance check and a unique reward ID, so copying a public transaction earns nothing and self-reported uptime does not count.

The first reward model will not pay "first to accuse" bounties. Even a true report could describe an incident the reward seeker created. Any future discovery program would need independent judging, tests against manufactured incidents, and a fixed budget.

How rewards are settled. Proposed steps: assigned, submitted, objectively checked, challenge window, decision. Duplicates resolve to one task, and nothing is paid while the evidence it relies on is provisional or unavailable. The question of who decides challenges must be answered before any real rewards exist; a majority opinion is not a substitute for an objective acceptance rule.

Penalties are narrow. Only provable, attributable protocol violations can be penalized, such as signing two statements the protocol forbids together, with the evidence, contest period and penalty set in advance. Disagreeing about a risk assessment, an AI mistake or a provider outage is not punishable fraud. Because deposits of a young coin may be worth little, deposits alone will not be relied on to secure the early network. Coin rewards also do not pay server bills, so TFA will publish a separate operating budget.

Why not proof-of-work. Bitcoin's proof-of-work secures its history by making rewrites costly. TFA is investigating a design without competitive puzzle-solving: useful, checkable work may determine rewards, while a separately specified consensus protocol secures TFA's ledger. The security cost of attacking a TFA chain has not yet been established.

The founder's coins. The founder may run nodes and earn under the same rules as everyone else. To limit concentration, founder nodes will be publicly labeled (labeling alone does not prove independence), and rewards will start only once outside operators participate, or early founder earnings will be capped.

TFA passes its own checks. Before launch, the coin must meet the standards TFA applies to others: published supply rules, a published threshold for what counts as a concentrated allocation, disclosed administrative powers, and no single party able to freeze holders. The official coin is defined by TFA's chain identity and genesis record, not by a Solana token address. The ticker TFA is also used by small unrelated tokens, so the official identifiers are the only reliable way to confirm the real coin. Practice coins run on a separate test network, may be reset, have no promised value, and will not convert into real TFA coins.

Still to be decided, only at the relevant stage: total supply and issuance, precision, budgets per role, fees, genesis allocation, founder treatment, staking, reward caps, upgrades and recovery, and who decides disputed tasks. The coin is independent of any other project's token, including WebbWing's.

Roadmap

TFA is at Stage 1, an offline prototype; each later stage starts only after the one before it is proven and approved. Rough time estimates are below; they are not promises.

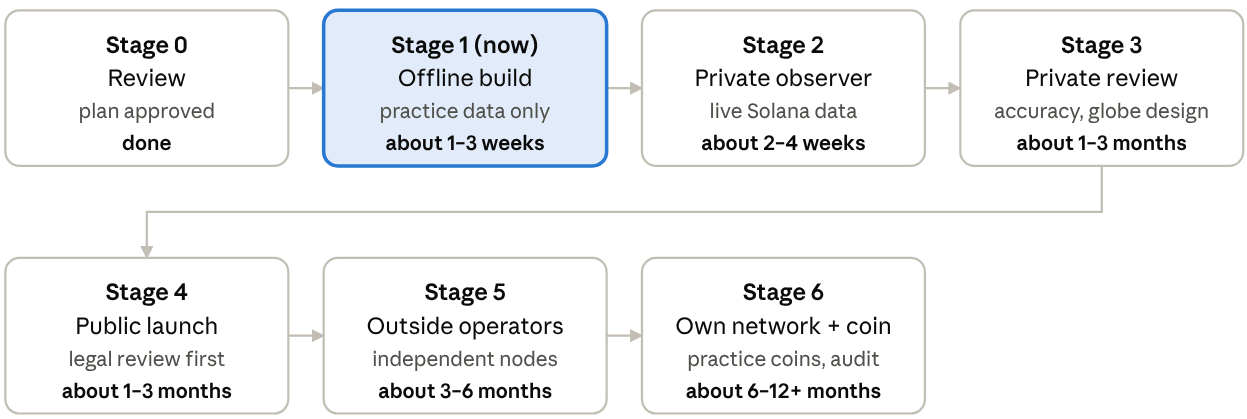

TFA roadmap · 7 stages, each needing its own approval

Each stage ends only when its evidence is in:

- **Stage 0 (done)** approved the plan, the threat model and the decision log.
- **Stage 1** builds the observer, the evidence ledger, money-trail tracing and launch linking on labeled practice data, including documented real rug pulls, honest launches and unclear cases. It ends with reproducible test runs, checked evidence hashes, passing positive and negative tests, and working corrections and missing-data cases. Measurement starts here, with a versioned test set kept separate from the data used to build the rules.
- **Stage 2** watches live Solana data privately, on a server separate from any trading system, with its own data keys and spending limits. It ends with named program coverage, restart and backfill tests, and tested handling of provisional data and source disagreement.
- **Stage 3** runs a held-out evaluation for each rule and protocol, measuring false alerts, missed events, delay and cost, restores a backup, and designs TFA's own globe for tfachain.com. Targets are set before the evaluation runs, and the globe design does not replace these tests.
- **Stage 4** publishes findings, the explorer and the globe only after legal review, an approved publication policy and a working dispute and correction process. The WebbWing globe may then read TFA's public data. Publishing findings does not mean the coin has launched.
- **Stage 5** brings in outside operators, measures their real independence, discloses shared dependencies, and tests conflicting sources and unavailable operators.

- **Stage 6** chooses the consensus and reward designs, runs them on an adversarial practice network, sets the genesis and supply policy, gets security and legal reviews, and then makes a separate, documented decision on launching the real coin.

Estimated timeline

As of September 2026, the estimates put private monitoring of live Solana data about 1–2 months away, a public launch at tfachain.com about 4–8 months away, and a real TFA coin about 1.5–2 years away if every stage goes well. The roadmap above shows the estimate for each stage.

The early stages are mostly building and move quickly. The later stages take longer because they depend on things that cannot be rushed: time watching real scams to measure accuracy, legal review, security audits, and recruiting independent operators. These estimates will be updated as each stage finishes.

Limits, risks and open questions

TFA is designed around its limits rather than hiding them.

Limits

- TFA sees only what its adapters support and its data sources provide. Unsupported programs, missing history and outages are shown as gaps.
- Privacy coins such as Monero are built to hide senders, receivers and amounts. TFA may see a supported transfer to an identified conversion service, but that does not give a complete trail into or out of a privacy system; a trace stops when the evidence cannot support the next link.
- A trace that reaches an exchange cannot identify the exchange's customer. TFA's globe shows relationships between addresses, not real-world locations: a wallet address carries no location.
- Until independent operators run nodes, TFA depends on its founder's servers and a small number of data providers.

Risks

- **False alarms** could harm innocent projects. Mitigations: private testing first, evidence on every report, careful wording, corrections, and disputes.
- **Legal exposure** from publishing reports and from issuing a coin. Mitigation: qualified legal review before either.

- **Gaming the rewards** through copied reports, fake incidents or many fake nodes. Mitigation: reward only hard-to-fake work, tested in simulation before real coins exist.
- **Conflicts of interest.** TFA applies the same rules to every project, including any linked to its founder or to sites that display its data, and logs every rule change.
- **Scammers adapt.** New tricks require new rules; TFA is never finished.
- **Misleading AI output.** Instructions hidden in token names or metadata could try to steer AI summaries. Mitigation: isolation, claim-by-claim evidence checks and human review before publication.

Open questions

- Who confirms that a report was right, for rewards and for public status?
- Coin supply, issuance, rewards per role, founder allocation, staking and penalties.
- How independent operators are admitted without letting one person pose as many.
- Which Solana launchpads and liquidity protocols are supported after the first ones.
- Who holds each kind of authority: protocol upgrades, rule approvals, emergency pauses and dispute decisions, and how those powers are recorded and transferred.

Disclaimer

This is a draft describing a project in early development. Nothing in it is an offer to sell anything, investment advice, or a promise that the TFA coin will exist, launch on any date, or have any value. Designs, numbers and defaults may change. TFA reports describe recorded blockchain activity and rule results; they are not determinations of wrongdoing by any person.