SaveSaveSaveSave

A Transparent Risk Tool for Everyone in Crypto

Verified. Constantly Improving.

> *“Four chains. Four data sources. Four reasons for the name: save your funds, save your time, save your trust, save the regret.”*

1. Executive Summary

Crypto is easy to get into and hard to get right. Anyone can send a transaction in thirty seconds. Almost no one gets a clear answer to the question that actually matters: how risky is what I'm about to do?

SaveSaveSaveSave answers that question. It checks a wallet or token against real security data — honeypot detection, mint authority, sanctions designations, malicious-address intelligence, and on-chain transfer patterns — and returns a plain verdict: **PASS, CAUTION, FAIL, or INSUFFICIENT DATA.** Not a price prediction. Not financial advice. A risk read, in seconds, before you act.

This edition documents a materially different product from the one the previous whitepaper described. What was a single-chain EVM tool is now a four-ecosystem screening engine — EVM (14 chains), Solana, Sui, and TRON — with wallet-risk intelligence that goes beyond a single vendor's opaque score: an independently verified on-chain sanctions oracle, malicious-address screening across two ecosystems, and a bounded, honestly labeled check for recent exposure to sanctioned counterparties.

Who it's for:

- People new to crypto, who shouldn't need to know what a “honeypot” or “mint authority” is before their first transaction is safe.
- Retail traders and small merchants, who need a fast, trustworthy check before interacting with something unfamiliar.
- Institutions and funds, who need a transparent, auditable risk layer — one they can inspect, not one they have to trust blindly.

2. The Name

A name built from one word repeated four times invites a fair question: why? The honest answer is that it's a deliberate structural choice, mapped onto four things the product actually does — not four abstractions invented to justify a word count.

- **Save your funds.** Catch honeypots, malicious contracts, and sanctioned or malicious-address exposure before a transaction is signed.
- **Save your time.** One scan, four blockchain ecosystems, and — where the chain supports it — multiple independent data sources queried in parallel.
- **Save your trust.** Every verdict shows its receipts. Each check states its source; a PASS is never presented as a safety guarantee.
- **Save the regret.** The Chainalysis Sanctions Oracle integration and recent counterparty checking exist to catch what a user would otherwise only discover after the money has already moved.

3. The Problem: Why Crypto Risk Is Still Invisible by Default

It's rarely the dramatic story. The common failure isn't "the market crashed" — it's smaller, quieter, and almost always preventable in hindsight: a token that turns out to be unsellable, built that way from the start; funds sent to an address that turns out, on public record, to be sanctioned; a wallet connected to a project without knowing what permission was just granted.

None of this requires bad luck. It requires the absence of one thing: a clear, honest signal *before* the decision, not an explanation after. Price charts tell you what already happened. Institutional compliance platforms are often genuinely good, but priced and sold through enterprise sales conversations — confirmed directly against how those vendors actually sell today, not assumed. There is very little built for the space between those two.

4. Architecture: What "Evidence-Based" Actually Means Here

- **Absence of evidence is never evidence of safety.** An incomplete check produces INSUFFICIENT DATA, not a quiet PASS.
- **A single confirmed critical finding overrides everything else.** Accumulated positive signals cannot outvote one verified critical one.
- **Every chain is evaluated on its own native security model** — not a shared template forced across ecosystems.
- **A PASS is a statement about coverage, not a guarantee.** Every result states how many checks returned usable data, and what confidence that supports.

This is enforced by a permanent, versioned regression suite — 61 automated tests covering malformed and partial provider responses, contradictory indicators, and adversarial cases written specifically to try to break the "weak evidence can't outvote strong evidence" rule.

5. Chain Coverage

EVM — 14 chains. Ethereum, BNB Smart Chain, Polygon, Arbitrum, Optimism, Base, Avalanche, Fantom, zkSync Era, Linea, Scroll, Blast, Mantle, Gnosis. Full token and wallet screening.

Solana. Its own native security model — mint/freeze authority, metadata mutability, transfer-hook presence — plus liquidity data shown as context, not folded into the score. Wallet screening reuses the same malicious-address data source as EVM, since address reputation isn't a chain-specific mechanic.

Sui. Token security only, scoped that way deliberately: the underlying data provider does not yet return holder or liquidity data for Sui. Its contract-function schema uses a value encoding distinct from every other chain here — both “1” and “2” signal “available,” caught by a purpose-built classifier.

TRON. Built on the same schema as EVM after confirming, not assuming, that TRON is served by the same general-purpose endpoint. A real disambiguation question with Solana's address format was resolved by calculation, not guesswork — roughly one in 400 quadrillion odds of collision for a genuine key.

Aptos — evaluated, not built. Named in the data provider's own marketing material; no dedicated technical endpoint could be confirmed against primary documentation, so it isn't claimed as supported.

Address type is detected automatically across all four ecosystems — the four formats (EVM's 0x-plus-40-hex, Solana's base58 without a leading zero, Sui's Move-type identifier containing ::, and TRON's T-prefixed Base58Check) are mutually exclusive by construction.

6. Wallet-Risk Intelligence: Beyond a Single Vendor's Score

Commercial blockchain-intelligence platforms are real and sophisticated, and in most cases entirely reasonable for what they're built to do. They are also, without exception in the direct experience of building this product, sold through enterprise sales conversations with no public self-serve pricing. The wallet-risk intelligence here was built around sources that are genuinely free, genuinely verifiable, and don't require a contract.

Layer 1 — Malicious-address screening

Sanctions exposure, phishing and wallet-draining association, theft and exploit involvement, money-laundering and mixer exposure, dark-web transaction history, cybercrime association. Available for EVM and Solana today.

Layer 2 — The Chainalysis Sanctions Oracle

A free, publicly callable smart contract, not a paid API. Confirmed independently: the Ethereum deployment address was cross-checked against Etherscan directly, and the exact function selector was computed with a hashing method sanity-tested against a universally known reference value first, then confirmed to appear, byte for byte, inside the actual deployed contract's bytecode. Eight of the nine chains where the oracle is deployed are supported here, three of which required a specific, non-obvious sub-network address rather than the chain's default one — caught by checking each directly, not by assuming a pattern.

One honestly disclosed limitation: an independent technical review found real gaps — in one case over ninety days — between an OFAC designation taking effect and the oracle reflecting it. A match is reported as exactly that: a match against this oracle, as of the check.

Layer 3 — Recent counterparty checking

Standard blockchain infrastructure has no method to retrieve “all transactions for an address” — that capability doesn't exist at the protocol level. What this checks instead: recent token-transfer activity, in a bounded, disclosed window (different real time spans per chain, since block speed varies enormously), cross-referenced against the same sanctions oracle. Native-asset transfers aren't visible to this method, and it says so. A single sanctioned counterparty moves the result to caution, not an automatic fail — that override is reserved for the wallet's own address being directly designated.

7. Transparency by Design

- Every check names its source — never “risky” without saying which provider or contract produced the finding.
- Coverage is stated, not implied: “17 of 18 checks completed” shown plainly.
- Confidence is separate from the verdict — full coverage reads differently from partial coverage, and the interface says so.
- Raw provider responses are one click away. Nothing is summarized in a way that can't be checked against its source.

8. Limitations, Stated Plainly

- **A clean result is not a safety guarantee.** It means no flags were found among what was checked — not that none exist.
- **Blockchain addresses are pseudonymous.** One entity may control many addresses; one address may be shared by many.
- **Labels can be wrong, and can change.** The Chainalysis oracle's own real-world staleness gap is evidence of exactly this lag.
- **Interaction with a flagged address is not proof of wrongdoing.** Exposure is reported as exposure, not guilt.
- **Coverage varies by chain, and that variance is disclosed** rather than presented as a false uniform experience.

9. What's Next

Two categories, kept honestly separate. Scoped and next in line: expanding counterparty checking's chain coverage as more infrastructure is verified, and reconsidering Aptos if a real data source is ever confirmed. Real but deliberately not started: continuous monitoring has a complete architectural design already written, held back because it requires this product to begin storing what a user is watching — the first genuine data-retention decision in its history, made on purpose rather than as a side effect of a feature.

10. Closing

Most tools in this space ask for trust. This one is built to need less of it — every verdict shows its sources, every limitation is stated instead of buried, and every technical claim in this document was verified against a primary source before being written down.

Four chains, several independently verified data sources, one evidence model applied consistently across all of them. **Verified**, because a security tool that asks to be trusted has an obligation to earn it. **Constantly improving**, because the risk landscape doesn't hold still, and neither should the tool built to read it.

This whitepaper is a living document. It will evolve as the product evolves.