

Know your risk before it knows you.

An AI-powered risk intelligence layer for DeFi
portfolios across Solana, Base, and Ethereum.

CONTENTS

Table of Contents

01	The Problem	03
02	The Solution	04
03	Position Detection Engine	05
04	The Risk Scoring Engine	07
05	AI Risk Intelligence	09
06	Architecture & Security	10
07	The RISK Token	11
08	Market & Competitive Landscape	13
09	Roadmap	14

ABOUT THIS DOCUMENT

This whitepaper describes RiskGuard's live product architecture and its planned token economy. The product described in sections 1–6 is operational today. The RISK token described in section 7 is a forward-looking design and has not been issued. See the disclaimer on the final page.

The Problem

Every DeFi investor knows where their money is. Almost none of them know how dangerous it is.

Decentralized finance has matured into a multi-chain landscape holding tens of billions of dollars in user assets. Yet the tooling available to the everyday investor has not kept pace with the risks. Portfolio trackers like DeBank, Zerion, and Zapper answer a single question well: *where are my assets?* They render balances, positions, and token lists across protocols. But they remain silent on the question that actually determines whether an investor keeps their money: **how much risk am I carrying, and what should I do about it?**

This gap is not academic. In 2025, DeFi protocols lost more than three billion dollars to exploits, hacks, and economic attacks — a figure that rose roughly forty percent year over year. The first quarter of 2025 alone was the worst three-month period in the industry's history. Behind each of these headline numbers are thousands of individual investors who held positions in affected protocols and had no warning, no risk signal, and no clear sense of their exposure until it was too late.

Three structural failures in current tooling

1. Balances are not risk

Knowing that a wallet holds a position in a lending protocol says nothing about the smart-contract risk of that protocol, the liquidity depth backing the position, the degree of admin-key centralization, or the market conditions that could unwind it. A portfolio that looks healthy by dollar value can be acutely fragile. Existing dashboards display the dollar value and stop there.

2. Solana is underserved

The dominant risk and analytics tools were built EVM-first. Solana — now one of the most active DeFi ecosystems — is either unsupported or treated as an afterthought. The protocols where Solana users actually deploy capital, such as Kamino, Jupiter, Raydium, and Orca, use position architectures that off-the-shelf trackers fail to detect at all. An investor cannot assess risk on positions a tool cannot even see.

3. Risk intelligence is locked behind institutional pricing

Where sophisticated risk analytics do exist, they are priced for funds and institutions. Platforms aimed at professional analysts begin at tens of dollars per month and climb into four figures, with interfaces designed for full-time analysts rather than the retail investor managing their own capital. The investor who most needs an accessible risk signal is the one least served by the market.

THE CORE INSIGHT

The DeFi market has comprehensive tools for *tracking* assets and institutional-grade tools for *analyzing* them — but nothing that gives the everyday multi-chain investor an honest, actionable risk signal across their entire portfolio, including the Solana positions other tools cannot detect.

The Solution

RiskGuard is a risk intelligence layer that sits on top of a user's DeFi portfolio and answers the question existing tools ignore: how dangerous is this, and what should I do?

A user connects their wallet — EVM or Solana — and within seconds RiskGuard detects every active position across supported protocols, scores the risk of each position along four independent dimensions, computes a consolidated portfolio risk score, and generates a plain-language analysis that explains the single biggest risk and recommends a concrete action.

The product is built on three technical pillars, each described in detail in the sections that follow:

Pillar one — Universal position detection

RiskGuard detects positions by calling smart contracts directly rather than relying on a hardcoded list of known receipt tokens. This allows it to surface positions that other trackers miss entirely, including Solana protocols that use program-derived addresses, NFT-based positions, native stake accounts, and proxy architectures. RiskGuard currently detects positions across more than twenty-five protocols spanning Base and Solana.

Pillar two — A transparent four-dimension risk engine

Every position is scored across smart-contract risk, liquidity risk, centralization risk, and market risk. These combine into a single portfolio score that also accounts for concentration using a Herfindahl-Hirschman Index measure of diversification. The scoring methodology is transparent and explainable — a user can see why a position scored the way it did.

Pillar three — AI-generated risk intelligence

Raw scores are translated into human language by an AI engine. Rather than presenting a number and leaving interpretation to the user, RiskGuard explains what is going well, identifies the single most pressing risk with specific figures, and recommends a concrete action — all in a few sentences.

25+ PROTOCOLS DETECTED	3 CHAINS COVERED	4 RISK DIMENSIONS
----------------------------------	----------------------------	-----------------------------

The result is a product that is accessible to a casual investor yet technically rigorous enough for an active DeFi trader — covering the chains and protocols that other tools leave in the dark.

Position Detection Engine

The hardest technical problem in portfolio risk is not scoring a position — it is finding it in the first place.

Most portfolio trackers detect positions using a simple heuristic: they look for known "receipt tokens" in the wallet. When a user deposits into a protocol, that protocol typically issues a token representing the deposit. If the tracker has that token's address hardcoded, the position appears. If it does not, the position is invisible. This approach fails silently and systematically for any protocol architecture that does not issue a simple, wallet-held receipt token.

RiskGuard takes a fundamentally different approach: it queries on-chain state directly, using detection methods tailored to each protocol's actual architecture. This is more engineering-intensive but produces reliable, comprehensive results.

Detection coverage

Base — 13 protocols

Protocol	Detection method
Aave V3	<code>balanceOf</code> on aTokens directly
Compound V3	<code>balanceOf</code> on Comet contracts
Uniswap V3	NonfungiblePositionManager (NFT)
PancakeSwap V3	NonfungiblePositionManager (NFT)
Aerodrome Slipstream	Direct NFTs + staked-gauge recovery via transfer history
Morpho Blue	ERC-4626 vault scan + transfer-history fallback
Moonwell	Strategy registry → proxy → balance
Steakhouse / Gauntlet	Curator-labeled Morpho vaults
Spark	ERC-4626 on sUSDC + sUSDS
Fluid	Dynamic factory enumeration + ERC-4626
Extra Finance (Lending)	LendingPool + dynamic eToken set
Extra Finance (Farming)	Leveraged vaults via event-log scan + struct decode

Solana — 15 protocols

Protocol	Detection method
Jupiter Lend	jITokens in wallet
Jupiter (Native Staking)	jupSOL mint + native stake accounts
Kamino Farms	PDA detection at known data offsets
Raydium CLMM	NFT positions via Helius DAS
Orca Whirlpool	Token + Token-2022 scan + batch metadata
Meteora DLMM	On-chain PositionV2 decode
MarginFi / Sentora	Program-account decode (fixed-point shares)
Save Finance (Solend)	Obligation accounts via program-account scan
Loopscale	Orderbook lender PDAs via program-account scan
Huma Finance	PST / mPST receipt tokens (USDC-backed)
Marinade Finance	mSOL mint + native stake accounts
Sanctum	INF mint address match
Jito	jitoSOL mint + staking detection
Hastra (RWA)	wYLDs + PRIME receipt mints

Coverage as of the document date and expanding. Native staking via Marinade and Jupiter is detected through Solana stake accounts, which issue no receipt token and are invisible to conventional trackers.

Architectures that hide positions

Through building this engine, RiskGuard has catalogued the specific architectural patterns that cause positions to be missed by conventional detection. Each represents a class of protocol that simple receipt-token scanning cannot handle:

- **Program-derived addresses (PDAs)** — protocols like Kamino Farms store positions in deterministic program accounts rather than wallet-held tokens, requiring direct account parsing at known data offsets.
- **NFT-based positions** — concentrated-liquidity protocols such as Raydium CLMM, Orca, and Uniswap V3 represent positions as non-fungible tokens, requiring digital-asset-standard queries and metadata resolution.

- **Token-2022 program** — newer Solana positions use the extended token program, which a scan limited to the original token program will miss.
- **Proxy and smart-wallet intermediaries** — aggregators route deposits through proxy contracts, so the underlying protocol position is held by a contract the user does not directly control.
- **Native stake accounts** — liquid-staking and native-staking products such as Marinade Native and Jupiter native staking issue no receipt token at all; the staked SOL lives in Solana stake accounts owned by the user, requiring direct stake-program queries to detect.
- **RPC enumeration limits** — some detection paths are blocked by RPC providers' limits on account enumeration, requiring transaction-history fallbacks.
- **Commercial name versus base protocol** — a position branded under one name in a third-party tracker may in fact be a wrapper around a different base protocol.

OPERATING PRINCIPLE

When a user reports a position that the detector returns as empty, RiskGuard treats this as a code problem, never a user problem. A mandatory diagnostic checklist is run before any position is concluded not to exist. This principle drives the engine's reliability and is the foundation of user trust.

The Risk Scoring Engine

A risk score is only useful if the investor understands what it measures and trusts how it was calculated.

RiskGuard scores every detected position across four independent dimensions, each capturing a distinct category of risk. The dimensions are weighted and combined into a single position score from 0 to 100, where higher indicates safer. Position scores then roll up into a consolidated portfolio score that additionally accounts for concentration.

The four dimensions

Dimension	Weight	What it measures
Smart Contract	30%	Contract-level security, audit status, and known vulnerability signals for the protocol holding the position.
Liquidity	25%	Depth of liquidity backing the position and the protocol's total value locked, indicating how easily a position can be exited.
Centralization	25%	Degree of admin-key control and governance centralization that could be used to alter or freeze the position.
Market	20%	The protocol's standing and the market conditions of the underlying assets, including volatility and impermanent-loss exposure.

The smart-contract dimension draws on dedicated security-data providers; the liquidity and market dimensions draw on protocol-level TVL and ranking data; and the centralization dimension incorporates admin-key analysis. Established, audited protocols are recognized as such and scored accordingly.

Portfolio-level scoring and concentration

Individual position scores do not tell the whole story. A portfolio with a single high-scoring position is still fragile because it is undiversified. RiskGuard therefore computes a concentration measure using the Herfindahl-Hirschman Index (HHI), a standard economic measure of concentration, applied to the distribution of capital across positions and protocols.

A well-diversified portfolio receives a diversification benefit; a concentrated one is penalized. An additional benefit is applied when a portfolio spans multiple chains, reflecting reduced single-ecosystem risk. The result is a portfolio score that reflects not just the quality of individual positions but the structural resilience of the whole.

WHY TRANSPARENCY MATTERS

Unlike opaque "risk ratings," RiskGuard's methodology is explainable. Each dimension maps to a clear, auditable data source, and a user can see which dimension is dragging a score down. This transparency is what turns a number into a decision an investor can act on with confidence.

AI Risk Intelligence

A score tells an investor that something is wrong. Intelligence tells them what to do about it.

The final layer of RiskGuard translates the quantitative output of the risk engine into clear, actionable language. This is powered by a large language model integrated directly into the analysis pipeline. Rather than presenting the user with a dashboard of numbers and leaving interpretation as an exercise, RiskGuard produces a concise, structured analysis for every portfolio.

The structure of an analysis

Each AI analysis follows a deliberate three-part structure designed to be read in under thirty seconds:

- **What is going well** — the strengths of the portfolio, grounding the analysis and avoiding pure alarmism.
- **The single biggest risk, with specific figures** — not a generic warning, but the one concrete exposure that matters most right now, named explicitly and quantified.
- **A concrete recommended action** — a specific, executable next step, such as reducing a particular allocation by a defined amount and redeploying it.

The analyses name protocols explicitly, cite real yield figures, and flag specific risks such as impermanent-loss exposure in a particular liquidity pair. Each carries a clear financial disclaimer: the analysis is informational and is not financial advice.

Efficiency and design

The AI layer is engineered for cost efficiency and speed. Analyses are generated using a fast, economical model and cached per session, keeping the marginal cost of each analysis to a fraction of a cent. This makes high-quality risk intelligence economically viable to offer broadly — including to free-tier users — rather than reserving it for premium customers.

Beyond static scoring — alerts and anomaly detection

Risk is not static. RiskGuard runs continuous monitoring that watches for material changes in a user's positions: sharp drops in protocol TVL, abnormal yield spikes, changes in admin-key configuration, and signs of liquidity exit. When a material risk event is detected, the user is alerted in the dashboard and, for high-severity events, in real time by email and optionally by Telegram, with a daily digest summarizing lower-priority signals. Users control which severities they receive and how often.

Performance history and exportable reports

RiskGuard periodically snapshots each user's portfolio value and renders a performance history chart that breaks down value over time by chain. Users can export their portfolio history as a CSV file and generate a complete portfolio risk report as a downloadable PDF — combining the consolidated risk score, the per-dimension breakdown, the full position table with range status, performance over the period, and the AI analysis into a single shareable document.

A NOTE ON YIELD REPORTING

For concentrated-liquidity positions, RiskGuard deliberately reports *range status* — whether a position is in-range, out-of-range, or full-range — rather than an estimated yield figure. An out-of-range position earns no fees and silently accumulates impermanent loss; surfacing that fact is more actionable than an approximate percentage. Fee-based yield is reported where it can be computed accurately.

THE PRODUCT THESIS

Detection finds the position. Scoring measures the risk. Intelligence tells the user what to do. Monitoring watches it over time. Together these four layers turn a portfolio of opaque positions into a continuously-assessed, clearly-explained risk picture.

Architecture & Security

RiskGuard is read-only by design. It never takes custody of funds and never requests permission to move them.

RiskGuard's relationship with a user's wallet is strictly observational. The application reads on-chain state to detect and score positions. It does not request token approvals, it does not initiate transactions, and it never holds user funds. Authentication is performed through message signing — Sign-In with Ethereum for EVM wallets and Ed25519 signature verification for Solana wallets — which proves wallet ownership without granting any spending permission.

Technical stack

Layer	Technology
Frontend	Next.js, React, TypeScript, Tailwind
EVM wallets	wagmi + viem + RainbowKit
Solana wallets	Solana Wallet Adapter (Wallet Standard)
Backend	Node.js + Express + TypeScript
Data	PostgreSQL + Redis
RPC & data providers	Helius (Solana), Alchemy (EVM), DeFiLlama, security-data APIs
Notifications	Email (Resend) + Telegram

Reliability engineering

Reading live blockchain state reliably at scale requires careful engineering. RiskGuard incorporates several reliability mechanisms developed in response to real-world failure modes:

- **Request coalescing** — concurrent requests for the same wallet are merged into a single in-flight computation, preventing redundant load and race conditions.
- **Non-decreasing confirmed cache** — a detection result is only allowed to replace a prior result if it accounts for at least as many positions, preventing transient RPC failures from causing positions to flicker out of view.
- **Automatic retry with backoff** — flaky RPC calls are retried with exponential backoff before falling back to the last known good result.

- **Transaction-history fallbacks** — where direct enumeration is limited by RPC providers, historical transaction analysis is used to recover positions.

SECURITY POSTURE

The combination of read-only access, signature-based authentication, and no custody of funds means that connecting to RiskGuard carries none of the smart-contract risk associated with granting token approvals. The worst-case data exposure is the public, on-chain contents of a wallet the user has chosen to connect.

The RISK Token

RISK is designed as a utility token that aligns the incentives of users, contributors, and the protocols RiskGuard analyzes — not as a speculative instrument.

IMPORTANT — FORWARD-LOOKING

The RISK token has not been issued. It is planned for a later phase, only after the product has an active user base and demonstrated revenue. Token design and timing are subject to change and to legal and regulatory review. Nothing in this section is an offer to sell or a solicitation to buy any token.

RiskGuard's product stands entirely on its own without a token. The RISK token is a mechanism layered on top of a working product to align incentives across the ecosystem and to fund continued development in a community-aligned way. Its design deliberately ties every unit of utility to real product usage.

Token utilities

1. **Subscription discounts** — paying for Pro or Whale plans in RISK confers a meaningful discount versus fiat pricing.
2. **Feature unlocking via staking** — staking RISK unlocks advanced analytics features without a recurring subscription.
3. **B2B API access** — protocols and applications consuming RiskGuard's risk API stake RISK in proportion to their call volume, tying network demand to token demand.
4. **Governance** — token holders vote on which new protocols and chains RiskGuard prioritizes for support.
5. **Contribution rewards** — users earn RISK for valuable contributions such as reporting false positives in detection or scoring, and through referrals.
6. **Priority analysis** — stakers receive more frequent risk re-analysis of their portfolios.

Distribution

RISK is planned as an SPL token on Solana with a fixed supply of one billion units. The distribution is weighted toward community and product, with founder and treasury allocations subject to extended cliffs and vesting to align long-term incentives.

Allocation	Share	Terms
Community airdrop	27%	Gradual unlock over 18 months
Product reserve	22%	Reserved for product development
Founders	20%	12-month cliff + 36-month linear vesting
Grants & partnerships	13%	Ecosystem development
Treasury / working capital	10%	6-month cliff + 36-month vesting
DEX liquidity	8%	At launch

Non-custodial staking

All staking is non-custodial and executed through audited smart contracts. RiskGuard reads stake status on-chain to verify entitlements but never takes custody of staked tokens. A third-party security audit is a prerequisite to launch. This architecture is consistent with the product's broader principle: RiskGuard reads the blockchain, it never holds user funds.

SEQUENCING

The token follows the product, not the other way around. The intended sequence is: mature product → active users → demonstrated revenue → ecosystem grants → token. This ordering protects both the product's credibility and the token's long-term integrity.

Market & Competitive Landscape

RiskGuard occupies a position no incumbent fully covers: accessible, multi-chain, risk-first, with native Solana detection.

The market for DeFi portfolio and risk tooling is established but structurally incomplete. Incumbents cluster at two poles: consumer trackers that show balances without risk, and institutional analytics priced and designed for professionals. RiskGuard targets the underserved middle — the self-directed investor who needs real risk intelligence at an accessible price, across the chains they actually use.

Competitive positioning

Product	Gap RiskGuard fills
DeBank	No Solana coverage, no risk scoring
Zerion	No risk analysis, no alerts
Nansen	Institutional pricing, complex, analyst-oriented
Zapper	EVM-only, no risk scoring
De.Fi	No complete portfolio view, no Solana

The complementary infrastructure layer

It is worth distinguishing RiskGuard from security-data infrastructure providers. Such providers operate at the infrastructure layer, selling risk APIs and SDKs to developers, wallets, and protocols. RiskGuard operates at the experience layer, building the consumer-facing product that turns raw risk data into decisions for the end investor. The two are complementary rather than competitive — RiskGuard consumes security data as one of several inputs to its scoring engine. RiskGuard is the experience layer that the infrastructure providers do not build.

DIFFERENTIATION IN ONE LINE

RiskGuard is the only product that covers Solana and Base natively, with a transparent risk score, AI-generated analysis, and position detection via direct contract calls — combining accessibility with technical depth.

Roadmap

RiskGuard follows a deliberate sequence: build a product people rely on, earn an active user base, then layer on monetization and community ownership.

Phase 1 — Foundation Complete

Monorepo infrastructure, automated deployment, multi-chain wallet connection for EVM and Solana, native balance reading, and consolidated portfolio display. Complete and operational.

Phase 2 — Core Product ~95% Complete

The position detection engine across more than twenty-five protocols on Base and Solana, the four-dimension risk scoring engine, and authentication via wallet signing. Operational, with systematic caching as the remaining refinement.

Phase 3 — Intelligence Substantially Complete

AI-generated risk analysis, the alerts and anomaly-detection system, email and Telegram notifications, per-position performance history with charting, CSV export, and exportable PDF risk reports are all live. This phase has progressed from partial to substantially complete, with ongoing refinement of yield accuracy across concentrated-liquidity protocols.

Phase 4 — Monetization & Growth Planned

Subscription plans (Free / Pro / Whale) with fiat and crypto payment, the B2B risk API, multi-wallet account support, and public launch. This phase establishes the revenue base that precedes any token consideration.

Funding strategy

RiskGuard pursues non-dilutive ecosystem funding in sequence with product milestones: community microgrants once a private beta is running, foundation-level grants once there is an active user base, and ecosystem hackathons once the product is monetizing with real traction. This funds development while preserving ownership and validating the product with the ecosystems it serves.

Program	Stage
Community microgrant	After private beta
Foundation grant	With active users
Ecosystem fund	In parallel
Hackathon & accelerator	After monetization is live

Disclaimer. This document is for informational purposes only and describes the current state and forward-looking plans of RiskGuard. The product features described in sections 1 through 6 reflect the operational product as of the document date. The RISK token described in section 7 is a forward-looking design that has not been implemented or issued; its design, terms, and timing are subject to change and to legal and regulatory review. Nothing in this document constitutes financial, investment, legal, or tax advice, nor an offer to sell or a solicitation of an offer to buy any token, security, or other instrument. RiskGuard's risk scores and AI-generated analyses are informational tools and not financial advice; users remain solely responsible for their own investment decisions. Forward-looking statements involve risks and uncertainties, and actual outcomes may differ materially. riskguard.finance