NightOvO

Solana Blockchain Intelligence Platform

Version 1.0

April 2026

nightovo.com

Confidential

Contents

01 — Abstract

NightOvO is a real-time blockchain intelligence platform purpose-built for the Solana ecosystem. After six months of research, prototyping, and iterative development beginning in October 2025, we are releasing this whitepaper to outline the platform's architecture, methodology, and vision.

The platform addresses a critical gap in the Solana tooling landscape: the absence of a comprehensive, accessible, and transparent scam detection system. Existing tools either provide surface-level analysis (RugCheck, Token Sniffer), cater exclusively to institutional budgets (Nansen, Arkham), or treat Solana as a secondary chain (GoPlus).

NightOvO combines twelve heuristic checks, graph-based deployer forensics, machine learning prediction, MEV detection, and social signal correlation into a single platform — delivered through a web dashboard, public API, Discord/Telegram bots, and a browser extension. All analysis is explainable: every risk score includes a complete breakdown of contributing factors.

> **Core thesis:** Security tooling for Solana should be transparent in methodology, accessible in pricing, and fast enough to warn users before they buy — not after.

02 — Problem Statement

The scale of token fraud on Solana

Solana's high throughput and low fees have made it the dominant chain for token launches. An average of 500–800 new tokens are deployed daily, the majority through automated launchpads like pump.fun. Our internal data collection (running since December 2025) indicates that approximately 45–55% of these tokens exhibit at least one strong indicator of fraudulent intent within the first 48 hours of trading.

The most common attack vectors, ranked by frequency in our dataset:

ATTACK VECTOR	PREVALENCE

		AVG. TIME TO EXECUTE
Liquidity removal (rug pull)	~38% of flagged tokens	2–24 hours
Mint authority exploitation	~25%	1–72 hours
Honeypot contracts	~18%	Immediate (by design)
Insider concentration with coordinated dump	~12%	4–48 hours
Name/metadata impersonation	~7%	Varies

Data collected from internal monitoring between Dec 2025 – Mar 2026. Sample size: ~42,000 token launches.

Limitations of current tooling

We evaluated every publicly available Solana token analysis tool during our initial research phase (Oct–Nov 2025). The findings shaped our design decisions:

TOOL	STRENGTH	CRITICAL GAP
RugCheck	Fast, free, widely adopted	No deployer history, no fund tracing, no alerting
GoPlus	Multi-chain coverage	EVM-first design; Solana checks are superficial
Token Sniffer	Contract analysis	No real-time monitoring, no graph analysis
Arkham Intelligence	Deep entity resolution	$500+/month; inaccessible to retail
Nansen	500M+ labeled addresses	$150+/month; limited Solana depth

03 — Solution Architecture

NightOvO is designed as a four-layer intelligence platform. Each layer operates independently but feeds into the layers above it, creating compound intelligence that improves over time.

FIGURE 1: PLATFORM ARCHITECTURE

L4	Distribution	Web Dashboard, REST API, Discord, Telegram, Browser Extension
L3	Intelligence	ML Prediction (XGBoost), Social Correlation, Entity Resolution
L2	Analysis	Risk Scoring, Deployer Forensics, Liquidity Intel, MEV Detection
L1	Ingestion	Solana RPC, Helius Webhooks, DEX Program Monitoring, Social Feeds

Data pipeline

The end-to-end latency from token mint to risk score availability is under five seconds in production conditions. The pipeline processes events asynchronously through Redis Streams with at-least-once delivery guarantees.

FIGURE 2: TOKEN ANALYSIS PIPELINE

1. Helius enhanced webhook fires on new SPL token mint event

↓

2. Event validated and queued to Redis Stream (`nightovo:analyze`)

3 Analyzer worker dequeues — fetches metadata, holders, deployer history via Helius DAS API

4 12 heuristic checks executed; risk score normalized to 0–100

5 Results persisted to PostgreSQL + TimescaleDB; graph data indexed to Neo4j

6 Live feed updated via SSE; if Danger (>75): alerts dispatched to configured channels

04 — Risk Scoring Methodology

Each token is evaluated against eleven independent heuristic checks, plus an optional ML prediction layer. Individual check scores are normalized to a 0–100 composite. The methodology was developed through analysis of approximately 12,000 confirmed rug pulls collected between November 2025 and March 2026.

FIGURE 3: SCORING CHECK WEIGHTS

Mint Authority Status	w: 15	Honeypot Simulation	w: 15
LP Lock/Burn Status	w: 15	Deployer Track Record	w: 15
Holder Concentration	w: 12	Freeze Authority	w: 10
Insider Clustering	w: 10	Hidden Fee Detection	w: 8
Name Impersonation	w: 8	Update Authority	w: 5
Supply Analysis	w: 5	ML Rug Probability	auxiliary

Safe	Caution	High Risk	Danger
0–25	26–50	51–75	76–100

Weights were calibrated iteratively against our labeled dataset. We prioritize recall over precision for the Danger classification — a false positive (flagging a legitimate token) is substantially less costly than a false negative (missing a scam). Current metrics on the held-out test set:

METRIC	HEURISTIC ONLY	HEURISTIC + ML
Precision (Danger class)	82.3%	87.1%
Recall (Danger class)	91.7%	94.2%

F1 Score	86.8%	90.5%

Evaluated on 3,200 labeled tokens (1,800 confirmed rugs, 1,400 legitimate). March 2026.

05 — Advanced Analysis Subsystems

5.1 Deployer forensics

Each deployer wallet receives a persistent reputation score (0–100) computed from their complete on-chain history. The scoring function considers: historical rug-to-launch ratio, wallet age, funding source tracing depth, and connections to flagged entity clusters. This score persists across all future launches from that wallet — serial scammers are identified before their next token appears on any DEX.

5.2 Fund flow graph engine

Wallet relationships are modeled in Neo4j as a directed property graph. Nodes represent wallets (labeled by type: exchange, bridge, deployer, whale, scammer); edges represent transfers (annotated with amount, token, timestamp, and transaction signature). The graph supports:

- Interactive visual exploration with incremental node expansion
- Shortest-path queries between any two wallets
- Community detection for entity resolution (Louvain algorithm)
- Ring and layering detection for obfuscation patterns

5.3 Liquidity intelligence

Continuous monitoring of liquidity pools across Raydium (AMM + CLMM), Orca (Whirlpools), Meteora (DLMM), and pump.fun (bonding curves). Key capabilities:

- **Wash trade detection** — Herfindahl index + round-trip scoring identifies volume inflation
- **Sniper bot detection** — flags addresses transacting within 3 slots of pool creation
- **LP status tracking** — lock duration, burn events, and removal alerts

5.4 MEV and sandwich attack detection

All swap transactions on monitored DEXes are parsed at the slot level. Sandwich attacks are identified by matching buy-swap-sell patterns from the same bot wallet

within a single slot on the same pool. Per-event data includes attacker profit, victim loss, and slippage impact.

5.5 Token similarity and impersonation detection

New tokens are compared against a registry of the top 500 legitimate Solana projects using: Levenshtein edit distance on names and symbols, Unicode confusable character detection (Cyrillic and Greek homoglyphs), and metadata fingerprinting. Similarity scores above 0.8 trigger automated alerts.

06 — Technical Infrastructure

COMPONENT	TECHNOLOGY	PURPOSE
API Server	Python 3.12, FastAPI	Async REST + WebSocket endpoints
Frontend	Next.js, TypeScript, Tailwind	Web dashboard with SSE live updates
Relational DB	PostgreSQL 16 + TimescaleDB	Token data, user state, time-series analytics
Graph DB	Neo4j 5	Wallet relationships, fund flow graphs
Message Queue	Redis Streams	Job distribution, pub/sub for real-time events
Blockchain Data	Helius API	Solana RPC, webhooks, parsed transactions
ML Inference	XGBoost, SHAP	Rug pull probability estimation
LLM	Ollama (self-hosted)	Natural language risk report generation
Containerization	Docker Compose (16 services)	Full stack orchestration

> **Self-hosted by design.** The entire platform runs on a single machine via Docker Compose. There is no dependency on cloud infrastructure, no vendor lock-in, and no third-party access to user data or query history. Operating cost is limited to the Helius API tier and domain registration.

07 — Business Model

Revenue is generated through freemium subscriptions and pay-per-use microtransactions, all settled on-chain in USDC via Helio (recurring) and Solana Pay (one-time).

Free

$0

- 10 scans/day
- 2-hop trace depth
- 3 watchlist wallets
- 7-day data retention

Pro

15 USDC/mo

- Unlimited scans
- Full trace depth
- 50 watchlist wallets
- AI risk narratives
- API access (1K/day)

Team

40 USDC/mo

- Everything in Pro
- 500 wallets
- Team sharing
- PDF report export
- API (10K/day)

Additional revenue from API tiers (25–90 USDC/month) and pay-per-use deep scans (0.50 USDC per scan). Monthly operating cost: $9–58. Break-even at 1–4 Pro subscribers.

08 — Development Timeline

Development began in October 2025 with initial market research and competitive analysis. The project has progressed through a structured seven-phase plan:

01	**Research**	Market analysis, competitor audit, Solana tooling evaluation	OCT 2025
02	**Data Collection**	Token launch monitoring, rug pull labeling, dataset construction	NOV–DEC 2025
03	**Foundation**	Core platform: scoring engine, Helius integration, API, live feed	JAN–FEB 2026
04	**Trace**	Neo4j graph engine, deployer forensics, holder analysis, d3 visualization	FEB–MAR 2026
05	**Intelligence**	Liquidity monitoring, whale tracking, MEV detection, copycat detection	MAR–APR 2026
06	**Distribution**	Discord/Telegram bots, browser extension, public API, auto-marketing	Q2 2026
07	**ML & Scale**	ML prediction model, lifecycle analytics, multi-chain expansion	Q3 2026

09 — Competitive Analysis

CAPABILITY	NIGHTOVO	RUGCHECK	GOPLUS	ARKHAM
Real-time scoring	✓	✓	✓	–
Deployer forensics	✓	–	–	✓
Fund flow tracing	✓	–	–	✓

MEV detection	✓	–	–	–
ML prediction	✓	–	–	–
Explainable scores	✓	Partial	Partial	N/A
Solana-first	✓	✓	–	–
Self-hostable	✓	–	–	–
Free tier	✓	✓	✓	–
Price (paid)	$15/mo	N/A	$99/mo	$500+/mo

10 — Conclusion

The Solana ecosystem generates enormous value — but it also generates enormous risk for participants who lack the tools to evaluate what they're buying. The tooling gap is not a technology problem; it's an access problem. The analysis capabilities exist, but they're locked behind institutional pricing, proprietary methodology, or incomplete implementations.

NightOvO is built on a straightforward premise: security intelligence should be transparent, explainable, and accessible. Every risk score comes with a complete breakdown. Every deployer profile shows the data it's built from. Every prediction includes the features that drove it.

We are not building a black box. We are building a flashlight.

OvO

nightovo.com · x.com/NightOv1215