



**NOXVIL**

PROJECT WHITEPAPER / PRODUCT & TECHNICAL DESIGN

# NOXVIL Whitepaper

Product and Technical Design / Version 0.1 / October 29, 2026

Six specialists. One inspectable workflow. Every decision has a reason.

---

## 01 / Executive Summary

NOXVIL is a collaborative terminal for trading research and strategy operations. It connects discovery, safety checks, position sizing, execution records, explanations and accounting in one inspectable workflow.

The current release is browser-based research software. The primary workspace uses deterministic synthetic quotes and a local ledger. A separate native workspace includes a read-only market adapter and research position tracking. Neither workspace signs transactions, connects a funded wallet or submits exchange orders. The named agents are rule-based software roles, with no LLM integration.

---

## 02 / Problem and Product Thesis

A recommendation alone does not explain how a decision became an action. NOXVIL retains the candidate, checks, parameters, action and accounting result together. Rejections are recorded with reasons.

The original overnight story of \$33.40 growing to \$13,456.22, 71 trades, 47 wins, 24 losses and 23 filtered tokens is an unverified concept scenario. It is not product performance, a backtest or an input to the current ledger. No return is promised.

---


## 03 / Product Principles

- Separate observation, safety, sizing, execution records and accounting.
  - Require explicit checks before creating a position.
  - Calculate exposure from available equity and configured limits.
  - Include adverse slippage and fees in research results.
  - Link decisions through trace IDs and a shared session.
  - Distinguish implemented functions, limitations and future plans.
- 

## 04 / System Architecture

The primary application combines the original terminal shell, a native right-hand dashboard, a deterministic state engine, a shared rendering controller and browser-local persistence. Shadow DOM isolates panel styling while both views consume the same session. The panel is an interactive interface rather than a screenshot or scaled iframe.

Build outputs include an online page and a self-contained offline page. Static panel content appears before runtime hydration. The separate native workspace is published under /native/. Its market access, model and view are independent modules. Static hosting does not provide a persistent trading backend.



---

## 05 / The Six Specialized Roles

| Role  | Responsibility                                      | Boundary                                   |
|-------|-----------------------------------------------------|--------------------------------------------|
| Scout | Discover pairs and volume candidates.               | Does not execute trades.                   |
| Guard | Check eligibility and reject failures with reasons. | A pass is not a comprehensive token audit. |

| Role    | Responsibility                                        | Boundary                                      |
|---------|-------------------------------------------------------|-----------------------------------------------|
| Sizer   | Calculate exposure from equity and allocation policy. | Cannot create funds or bypass limits.         |
| Trigger | Record research entries and exits.                    | No signing or real order submission.          |
| Scribe  | Preserve explanations under a trace ID.               | Local records are not independently attested. |
| Ledger  | Reconcile settlement, cash and equity.                | Research balances are not custody balances.   |

The operator reviews policy, interprets data quality and exports records. Role names describe responsibilities, not independent intelligence or security assurance.

---

## 06 / Data and Operating Modes

The primary workspace uses deterministic synthetic fixtures. They reproduce discovery, rejection, sizing, entry, exit and settlement without a funded account. They are not live prices.

The independent workspace contains a read-only DEX Screener adapter. Access depends on upstream availability, the network and browser or proxy configuration. Failed access produces a data-state message, without fabricated live-price fallback. Market connectivity was unavailable in the validation environment.

Published prices and observed liquidity do not guarantee executable fills. Available market checks do not constitute complete contract analysis or a honeypot audit.

---

## 07 / Decision Lifecycle

Scout proposes a candidate. Guard records a pass or rejection. Rejections leave a trace without increasing the trade count. Sizer determines allocation for an eligible candidate. Trigger records an entry and an open position, followed by a later close. Scribe preserves reasons, and Ledger reconciles the settled result.

The primary engine intentionally exposes the open state before settlement. The independent research engine monitors configured exits when data is available. Closing the browser stops browser execution; no continuously running overnight backend is included.

---

## 08 / Risk Policy and Defaults

These are independent research workspace defaults, not proven trading recommendations. The primary deterministic workspace has its own fixture checks.

| Parameter                         | Default or boundary                                  |
|-----------------------------------|------------------------------------------------------|
| Initial research equity           | \$33.40                                              |
| Minimum observed liquidity        | \$25,000                                             |
| Minimum one-hour volume           | \$1,000                                              |
| Selling activity                  | Recent sells required by available-data checks       |
| Maximum fetched-data age          | 45 seconds                                           |
| Allocation                        | 10% of available equity; configurable from 1% to 25% |
| Slippage or indicative impact cap | 1.5%                                                 |
| Modeled fee                       | 0.1% per side                                        |
| Take-profit threshold             | 4%                                                   |
| Stop-loss threshold               | 2%                                                   |
| Maximum holding time              | 180 seconds                                          |
| Session loss stop                 | 10%                                                  |

Liquidity and volume can be misleading or manipulated. Research exit thresholds do not guarantee achievable exit prices. Production operation requires additional market, contract, custody and execution controls.

---

## 09 / Costs and Ledger Invariants

Research accounting includes configured adverse slippage and per-side fees. A simplified long-position model is:

```
entryFill = observedEntry * (1 + entrySlippage)
quantity = allocatedCash / (entryFill * (1 + entryFee))
exitFill = observedExit * (1 - exitSlippage)
netProceeds = quantity * exitFill * (1 - exitFee)
realizedPnL = netProceeds - allocatedCash
settledEquity = initialEquity + sum(realizedPnL)
```

Exact arithmetic and rounding follow the relevant engine. The two workspaces have separate session stores and accounting models; their results are not a single verified performance series.

Only closed trades affect realized P&L. Rejections and open positions do not count as realized gains. Available cash deducts committed position costs. Mark-to-market equity, settled balance and spendable wallet funds are distinct concepts.

Core invariants cover finite nonnegative balances, valid configurations, position transitions and reconciliation of initial equity plus realized results. The curve represents recorded research settlement points.

---

## 10 / Records and Persistence

Trace IDs connect candidate checks, sizing, execution and settlement. The interface supports role filters, trace inspection, JSON session exports and CSV trade exports.

Sessions use browser-local storage. Records and pending research positions can be restored after refresh, but storage can be unavailable, cleared or manually edited. Records are not signed, tamper-proof or independently audited. Export important records regularly.

The default primary cadence is 4x, approximately 0.75 seconds per workflow step. This is a replay presentation setting, not market latency or trading throughput. Saved preferences can change later visits.

---

## 11 / Security and Operational Boundaries

The release does not custody assets, request seed phrases, sign transactions or submit orders. No LLM credentials or autonomous model execution are part of the current engine.

Production extensions would require isolated signing, permission limits, idempotent submission, transaction confirmation, reconciliation, duplicate-order protection, exposure limits, emergency controls and monitored recovery. External data requires freshness checks, validation and explicit

failure handling.

Local research balances are not proof of funds, real execution or investment performance.

---

## 12 / Roadmap

### Implemented

- Original terminal presentation and native dashboard.
- Six rule-based roles and a shared research session.
- Visible entry and settlement stages with configurable limits.
- Persistence, trace inspection and exports.
- Independent read-only market adapter and research ledger.

### Research and Evaluation

Improve data resilience, safety evidence, historical datasets and evaluation baselines. Assess future strategy intelligence independently from interface animation and ledger mechanics.

### Persistent Services

Develop monitored backend operation with durable records, scheduling, recovery and observability. Browser operation must not be confused with persistent services.

### Constrained Execution

Explore permission-limited execution only after signing isolation, reconciliation, risk controls and failure recovery are implemented and reviewed. These capabilities are future work, without a guaranteed delivery date.

---

## 13 / Token and Ecosystem Position

This release defines no issued token, allocation schedule, fundraising arrangement, staking program, yield entitlement or revenue distribution. CA: TBA is an address placeholder, not confirmation that a contract exists or will be deployed.

GitHub is the source reference. A social button without a configured account is a reserved entry. Any future token proposal requires separate verifiable documentation.

---

## 14 / Validation

Checks focus on sizing limits, visible open-to-close transitions, costs, ledger reconciliation, filters, exports, persistence and responsive layouts. Browser checks confirm shared state, native panel rendering and usable layouts. Document checks cover HTML and PDF delivery.

Passing software checks does not establish profitability, token safety or production trading readiness. Data availability and execution require separate evidence.

---

## 15 / Limitations

- Synthetic quotes cannot establish a strategy edge.
- Indicative market data does not guarantee executable liquidity.
- Closing the browser stops execution.
- Local records can be changed or lost.
- Partial checks cannot certify token safety.
- Modeled fees do not include every real-world cost, network fee or execution effect.
- Research results are not offers, guarantees or forecasts.

NOXVIL makes the decision chain inspectable. It does not eliminate market risk or replace a tested production trading system.

---

## 16 / Sources and Version History

Project source: <https://github.com/GHOSTXrun/>-

Project workspace: <https://ghostxrun.github.io/-/>

Market adapter reference: <https://docs.dexscreener.com/api/reference>

Version 0.1 / October 29, 2026 documents the current terminal, rules, research modes, accounting, boundaries and roadmap. This English edition accompanies the English frontend. Future implementation changes should update this document and remain traceable in the repository.

---

NOXVIL · v0.1 · 2026-10-29 · Product and implementation boundaries are versioned with the source.