

PLANCKNET

Production-Grade Execution Kernel

PRE-SEED INVESTMENT MEMORANDUM

Raising \$200,000

SAFE • \$2.5M – \$3.0M Valuation Cap • 20% Discount

December 2025

CONFIDENTIAL

Executive Summary

"We've built the deterministic risk simulation kernel institutions wish they had before the last crisis."

PlanckNet is a **production-grade deterministic execution kernel** for institutional risk simulation. Unlike traditional tools that aggregate risk into opaque metrics, PlanckNet enables controlled, repeatable stress experiments that decompose portfolio losses into their fundamental components: market risk, liquidity execution cost, and stabilization controls.

Our platform transforms stress testing from a compliance checkbox into an experimental discipline, allowing hedge funds, asset managers, and banks to understand exactly *why* portfolios fail under stress—and how to prevent it.

Investment Highlights

Stage	Pre-Seed
Raising	\$200,000
Instrument	SAFE with \$2.5M-\$3.0M cap, 20% discount
Team	3 founders (Quant, Engineering, Business)
Product	Production-grade kernel with live API
Target Market	Hedge Funds, Asset Managers, Banks

The Problem

Institutional investors don't understand why portfolios break under stress. This isn't a minor inconvenience—it's a fundamental failure in risk management that has repeatedly amplified losses during major crises.

Current Pain Points

- **Black-Box Risk Metrics:** Standard tools aggregate risk into single numbers (VaR, ES) that hide failure modes
- **Non-Reproducible Tests:** Stress tests produce different results on different runs, making debugging impossible
- **Mixed Components:** Liquidity impact, shock propagation, and stabilizing rules are tangled together
- **Compliance-Driven:** Current stress tests are checkbox exercises, not analytical tools

The result: When a portfolio loses money in a crisis, institutions cannot determine whether the loss came from market moves, liquidity costs, or control rule failures. This makes prevention and improvement impossible.

Why Existing Solutions Fail

Solution	Approach	Limitation
Backtesting Tools	Historical replay	Non-experimental, past-only
Compliance Tests	Regulatory scenarios	Checkbox-driven, not analytical
Internal Quant	Custom simulators	Often exceeds \$2M/yr, opaque
PlanckNet	Deterministic experiments	Full decomposition & replay

The fundamental gap: None of these solutions allow controlled, repeatable stress experiments with decomposed PnL attribution. They tell you what happened, not why.

Our Solution

PlanckNet is a production-grade deterministic execution kernel that provides institutional-grade experimental risk analysis. Our platform enables portfolio managers to run controlled stress experiments with full reproducibility and granular PnL decomposition.

Core Architecture

- **Risk Propagation Operator (RPO):** Models how factor shocks propagate through correlated instruments
- **Liquidity Shock Operator (LSO):** Quantifies execution costs and market impact under varying liquidity regimes
- **Drift Guard System:** Per-portfolio circuit breakers that can be tested before deployment
- **Deterministic Execution Kernel:** Seeded RNG ensures identical runs produce identical results for debugging

Key Differentiator: Determinism

In real institutions, risk decisions must be explainable, stress results must be reproducible, and controls must be tested—not assumed. Stochastic backtests hide failure modes. Deterministic simulation exposes them. This transforms stress testing into an experimental discipline.

Technology Deep Dive

PlanckNet V18.3.4.2 represents over 1,700 lines of formally specified execution logic, implementing a multi-portfolio simulation kernel with mathematically rigorous operators.

What "Production-Grade" Means

- **Deterministic replay tests:** CI validates reproducibility across runs
- **Versioned specification:** Formal spec with semantic versioning
- **Structured artifacts:** Every run produces auditable JSON output
- **Containerized deployment:** Docker-ready for VPC/on-prem

Technical Specifications

Kernel Version	V18.3.4.2-FINTECH
Specification	1,700+ lines formal pseudocode
Core Operators	RPO, LSO, Drift Guard
Regime Support	Normal/Elevated/Extreme volatility; Normal/Mild/Severe liquidity
API	FastAPI REST with bearer auth
Compute	PyTorch tensor operations (CPU/GPU)

PnL Decomposition

Every simulation step provides granular attribution:

- **Market PnL:** Profit/loss from price movements on existing positions
- **LSO Cost:** Liquidity execution cost from trading activity
- **Raw PnL:** Combined market + liquidity before controls
- **Net PnL:** Final PnL after drift guard intervention
- **Avoided Loss:** Quantified protection provided by circuit breakers

Use Cases

PlanckNet serves as pre-trade risk intelligence—helping institutions understand portfolio behavior before crises occur.

Primary Applications

- 1. Strategy Comparison:** Compare portfolio strategies under identical stress regimes
- 2. Liquidity Attribution:** Quantify what portion of tail losses come from liquidity costs versus market moves
- 3. Control Validation:** Test circuit breakers and risk guards before deployment, not during crises
- 4. Tail Risk Analysis:** Understand recovery dynamics and path-dependent behavior under extreme scenarios
- 5. Regulatory Stress Testing:** Generate reproducible, audit-friendly stress test results

Market Opportunity

Why Now

- **Regulatory Pressure:** Post-2023 banking crisis, institutions increasingly need explainability and reproducibility for governance
- **Liquidity Focus:** SVB and Credit Suisse failures highlighted liquidity risk as a primary concern
- **Portfolio Complexity:** Multi-asset, cross-border portfolios require more sophisticated simulation tools
- **Infrastructure Cost:** Internal quant infrastructure can exceed \$2M+ annually and is often brittle

Target Market

Segment	Buyer	Sales Cycle	Priority
Hedge Funds	CIO / Risk Head	3-6 months	Primary
Asset Managers	Risk Committee	6-9 months	Primary
Banks	CRO / Model Risk	9-12 months	Secondary

Business Model

Enterprise SaaS / Research Platform with annual contracts and pilot-to-deployment conversion.

Pricing Structure

Phase	Duration	Price
Paid Pilot	90 days	\$5,000-\$10,000/month
Annual License	12 months	\$100,000-\$150,000 ARR
Enterprise	Multi-year	\$200,000+ ARR

Sales Motion

We employ a hypothesis-driven pilot approach:

1. Initial diagnostic call to understand current stress testing process
2. Define 2-3 testable hypotheses (e.g., "X% of tail losses come from liquidity")
3. 90-day paid pilot with structured deliverables
4. Success measured by decision-making impact, not feature usage
5. Convert to annual license with expansion path

Current Status & Traction

Product Development

- ✓ **Production-grade simulation kernel** (V18.3.4.2) — not mockup
- ✓ **Live API + console demo** — integration-ready REST API
- ✓ **Formal specification** — 1,700+ lines PyTorch implementation
- ✓ **Multi-portfolio support** — three archetype strategies validated
- ✓ **Deterministic reproducibility** — validated across multiple runs

Validation & Pipeline

- Academic advisor — methodology validation in progress
- Design partner pipeline — risk team conversations active
- Institutional use cases — defined based on market research

Team

We build research-grade systems that can ship. Our team combines deep quantitative expertise with production engineering discipline.

Founding Team

- **Quant / Mathematics Lead:** Risk modeling, simulation logic, mathematical foundations
- **Engineering Lead:** Deterministic execution, infrastructure, API design
- **Business Lead:** Institutional fintech experience, go-to-market strategy

Advisory

Academic advisor for methodology validation and research credibility (in progress)

The Ask

We are raising **\$200,000 pre-seed** to deploy our validated kernel into initial institutional pilots and harden the API infrastructure.

Amount	\$200,000
Instrument	SAFE (Simple Agreement for Future Equity)
Valuation Cap	\$2.5M - \$3.0M
Discount	20%
Implied Dilution	~6-8%

Use of Funds (9 Months)

Category	Amount	%
Team (3 founders × 9 months)	\$100,000	50%
Pilot Deployment & Sales	\$50,000	25%
Infrastructure (compute, data, legal)	\$30,000	15%
Buffer / Contingency	\$20,000	10%
TOTAL	\$200,000	100%

9-Month Milestones

This round funds the path from validated kernel to institutional pilots and Seed-readiness.

Month	Milestone
1-3	Complete validation suite; API hardening; first design partner signed
4-6	Engineering whitepaper; 2-3 paid pilots underway
7-9	First ARR; Seed round preparation with case studies

Path to Seed Round

Successful completion of pre-seed milestones positions us for a \$1-1.5M Seed round at 12-18 months, targeting:

- 3-5 paying institutional clients
- \$300K-500K ARR
- Validated sales motion and ICP
- Research publication or whitepaper

Competition & Moat

Competitive Landscape

We are not replacing Bloomberg or MSCI. We are the **experimental layer** that runs alongside existing tools.

- **vs Vendor Suites (MSCI, Bloomberg):** Monolithic suites that typically don't provide deterministic replay + stepwise decomposition as a first-class workflow.
- **vs Internal Quant Infra:** Buy beats build. Internal simulators cost \$2M+/yr and are often brittle.
- **vs Monte Carlo:** MC gives distributions; we provide deterministic replay and attribution for debugging.

Our Moat

- 1. Deterministic Execution:** Same seed = same crash path = debuggable
- 2. Decomposed Attribution:** Market vs liquidity vs controls — not aggregated
- 3. Formal Specification:** 1,700+ lines of auditable logic
- 4. Verification Suite:** Reproducibility validated across runs

Architectural moat: Incumbents sell stochastic suites. Retrofitting determinism into existing engines is hard. Our advantage compounds with each pilot deployment — validated scenario packs, calibration playbooks, and institutional learnings create defensible network effects.

Vision

"The standard experimental risk simulation layer for institutional portfolios."

We envision a world where:

- **Risk is decomposed, not aggregated** — every loss has a traceable cause
- **Stress tests are experiments** — reproducible, debuggable, improvable
- **Controls are tested, not trusted** — circuit breakers validated before crises

Think of PlanckNet as **Docker for risk strategies** — a standardized experimental layer for shipping and testing risk controls. Reproducible stress packs, audit-friendly, vendor-neutral. That's the platform play.

Contact

Live Demo Available Upon Request

founders@plancknet.io

Thank you for your consideration.