

Crunch × LysLab LowFreq Alpha Engine

The Proposition

CrunchDAO brings 11,000 ML engineers. LYS Labs brings the fastest on-chain data and execution infrastructure on Solana. Together, we create the first crowd-sourced alpha discovery platform with institutional-grade, sub-300ms signal-to-execution. No quant competition platform in crypto can offer this full-stack latency profile today.

End-to-End Pipeline Performance

The LYS Flash stack delivers a complete signal-to-mempool pipeline in approximately 236ms, broken down across three layers:

< 100µs TX Build (Rust SDK) 523ns instruction extraction • 915x faster than Node.js	< 10ms TEE Signing VSOCK local comms • 20-50x faster than 3rd party	225ms P50 Transport to Mempool #1 ranked • 100% success rate • 0% drift
--	--	---

Benchmarked on Solana mainnet: 600 transactions across 100 iterations against 6 transport layers, including Nozomi, Helius, Jito, and vanilla RPC. New benchmarks will be released on Feb 19th, which we expect will be even better than the existing ones.

Transport Layer Benchmark: LYS Flash vs. Competitors

Transport	P50 Confirm	Slot 0-1 Rate	Success
LYS Flash	566.6ms	98%	100%
Nozomi (market leader)	566.3ms	97%	100%
Zero Slot	580.3ms	98%	100%
Helius	586.0ms	96%	100%
Vanilla RPC	1,126.3ms	31.6%	94%
Jito	1,256.8ms	26%	100%

LYS Flash achieves 2x faster confirmations than vanilla RPC and 2.2x faster than Jito, with the tightest slot placement precision (only 2% landing in Slot +2, the lowest outlier rate across all transports). For CrunchDAO quants, this means discovered alpha signals are executed before they decay.

TEE Wallet Management: The CrunchDAO Multiplier

The TEE Wallet Service solves a critical scaling challenge for the partnership: running hundreds of quant strategies simultaneously with isolated, secure wallets.

- **Hardware-Isolated Security:** AMD SEV-SNP enclaves ensure private keys never leave the secure environment. Even cloud admins cannot extract them.
- **Infinite Scalability:** Deterministic key derivation from a single Master Key means each CrunchDAO submitter gets a unique wallet with zero storage overhead. No database of millions of private keys.
- **Multi-Party Control:** Shamir's Secret Sharing splits the master key across trusted parties (e.g., 3-of-5 threshold), eliminating single points of failure for pooled capital.
- **Sub-10ms Signing:** VSOCK local communication eliminates the 200-500ms penalty of third-party wallet providers. This is the answer to Jean's open question on TEE speed — it adds negligible latency to the pipeline.

Implementation Pathways Mapped to the Stack

Jean's scoping document outlined three submission types. Here's how each maps to LYS infrastructure and recommended prioritization:

Pathway	Description	Latency Need	Priority
Probabilistic Modeling	ML engineers submit return distributions or Gaussian mixtures. Meta-model aggregates. LYS Flash executes.	Low (submission) / High (execution)	Phase 1 ★
Directional Probability	Up/Down signals with confidence %. Scoring on seconds/minutes cycle.	Medium	Phase 1 ★
Order Level	Tick-by-tick trade submissions. Requires full ~236ms pipeline per tick.	Critical	Phase 2

Recommendation: Launch with Probabilistic Modeling and Directional Probability in Phase 1. These are the most ML-native formats, will engage the largest share of CrunchDAO's 11k engineers, and don't require tick-level latency for submissions. The meta-model output then feeds into LYS Flash for institutional-speed execution. Graduate to Order Level in Phase 2 once the pipeline is validated.

The Data Moat: Enigma Condition Engine

Enigma's condition builder provides 50+ parameters across 10 categories that form the feature space for ML model training. This is on-chain behavioral microstructure data that simply does not exist in traditional quant competitions:

Category	Key Signals for ML
Price Metrics	Current price, ATH, local ATL, market cap, price-to-ATH distance
Volume Metrics	Buy/sell counts, buy/sell volume, pro trader buys
Holder Dynamics	New vs. aged wallet %, top holder concentration, dev holding %
Developer Activity	Dev has sold, partial sells, buys before dev sell, dev active
Swap Events	Gross amounts, DEX fees, creator fees, pool liquidity, wallet balance Δ
Graduation Events	Migration type, graduation liquidity
Technical Indicators	RSI, MACD, Bollinger Bands, volatility, momentum reversal
P&L Conditions	Profit targets, stop loss, trailing stop, price drop from peak

The unique edge: Holder dynamics and developer activity data (wallet age distribution, dev selling patterns, new wallet inflows) are signals that Numerai and other traditional quant platforms cannot offer. These behavioral features are where CrunchDAO engineers will discover genuinely novel alpha.

Data Obfuscation Strategy

To ensure interdependency between CrunchDAO submitters and LYS Labs (preventing quants from front-running specific tokens independently):

- **Token identifiers stripped:** Assets referenced by anonymous IDs, not mint addresses.
- **Time-shifted windows:** Historical data chunks offset by variable delays to prevent real-time replication.
- **Feature normalization:** Absolute values converted to z-scores or percentile ranks where possible, preserving statistical relationships while obscuring raw data.

This creates the right incentive structure: CrunchDAO engineers build generalizable models on rich features, but need LYS Flash for live execution on identified assets.

Revenue Architecture

Three-layer monetization that aligns incentives across both parties:

1. **Competition Layer (CrunchDAO):** Platform fees on competition entry, data access subscriptions, and model submission infrastructure. CrunchDAO's existing monetization model applies directly.
2. **Execution Layer (LYS Labs):** Volume-based fees on all trades executed through LYS Flash. As winning models generate more trading activity, LYS captures proportional revenue.
3. **Capital Ramp (Joint):** Structured funding ladder for top performers (\$5k → \$10k → \$30k+). Performance-based capital allocation from a pooled fund, with returns split between the quant, CrunchDAO, and LYS Labs.

Stretch goal — Community Market Maker: An ensemble meta-model aggregating hundreds of probabilistic submissions to provide non-adversarial liquidity. LYS Flash's 100% reliability and zero drift beyond 2 slots makes this operationally viable. No single participant controls the model; only the meta-model executes.

Responses to Open Scoping Questions

"Do we need to colocate?"

Not necessarily. The TEE wallet service communicates via VSOCK (local virtual socket), achieving sub-10ms signing without geographic colocation. The transport layer benchmark was run on Solana mainnet under production conditions and achieved 225ms P50 without dedicated colocation. Colocation becomes relevant only if we pursue the Order Level pathway at tick granularity, and even then only for marginal improvements.

"TEE speed?"

Sub-10ms for transaction signing via AMD SEV-SNP enclaves with VSOCK communication. This adds negligible overhead to the ~236ms total pipeline. The TEE also provides deterministic wallet derivation, meaning we can spin up isolated wallets for each CrunchDAO submitter at zero storage cost.

"Beeks or interconnect?"

LYS Flash's custom transport layer already outperforms or matches all tested alternatives including Nozomi (the current fastest public option). Beeks interconnect may provide marginal gains for the Order Level pathway but is not required for Phase 1. Recommend evaluating after Phase 1 results.

"Asset universe?"

Recommend starting with graduated tokens on Raydium. The graduation event is already captured in Enigma's condition engine (migration type + graduation liquidity), these tokens have

sufficient trading volume for meaningful execution, and the graduation signal itself is a rich feature for ML models to learn from.

Proposed Next Steps

1. **Technical Requirements Table** (Philippe + LYS CTO): Define data obfuscation specs, API handoff protocols, and sandbox environment configuration for the local runner.
2. **Phase 1 Roadmap** (Andra + Jean): Timeline for Probabilistic Modeling and Directional Probability competitions. Target launch window, scoring criteria, and leaderboard mechanics.
3. **Economics & Revenue Split** (Andra + Jean): Formalize the three-layer revenue model, capital ramp thresholds, and Community Market Maker governance structure.
4. **Online Runner Integration** (Philippe + LYS CTO): Define the integration between CrunchDAO's Mid+One runner and LYS Flash's execution pipeline. TEE wallet provisioning workflow for submitted models.