INVESTOR BRIEF · v0.1 · 2026-05-20

SPREADR

The institutional-grade autospreader for perp DEX arbitrage.

Spreadr executes net-funding arbitrage and basis trades across five perpetual DEXes — Hyperliquid, Extended, Lighter, Orderly, and Aster — with milli-second leg matching, strict risk controls, and a full audit trail. It runs co-located in Tokyo for low-latency access to the major venue endpoints, and isolates each user's trades with per-spread runtime + per-user-encrypted credentials.

CATEGORY
Perp DEX execution infrastructure

STAGE
Early-access cohort

FOUNDED
2026 · hej, Tokyo

CONTACT
spreadr.xyz

Executive summary

Perpetual decentralized exchanges (perp DEXes) settled **over $2 trillion in volume in 2025**, with the top five venues now operating 24/7 markets that match the depth and tightness of major centralized exchanges. The market has matured fast — but the tooling to execute **across** venues has not. Most cross-venue spread capture today is done by manual desks, proprietary HFT shops, or bespoke scripts.

Spreadr productizes that capability. It is the first institutional-grade automated cross-DEX-perp spread trader: pick two venues, set a spread, set a size — Spreadr executes both legs in lockstep, hedges every fill before it can drift, and runs the trade end-to-end with no human supervision.

The product in one paragraph

A user connects their wallets to Spreadr (Privy auth), adds API keys for each exchange they want to trade on (encrypted under a user-scoped context, never readable by Spreadr operators), then configures a spread trade: market, two venues, target spread, target size, risk caps. Spreadr spawns an isolated worker for that trade, connects to both venues directly, posts a maker order on one side, and slams a taker hedge on the other the moment the maker fills. Every fill is durably logged before its hedge is sent, so no crash can leave the user one-legged. Trades can run for minutes to days; users can amend, pause, or cancel at any time.

Why it works

- **Latency.** Tokyo-colocated, single-digit ms to most exchange endpoints.
- **Isolation.** Per-trade runtime with its own network egress — trades never compete for the same rate-limit budget.
- **Custody.** Credentials are encrypted under a user-scoped context. No operator at Spreadr can decrypt user keys, by construction.
- **Reliability.** Every fill goes to a durable ledger before its hedge is sent. Worker crash → recovery worker reconciles, no double-hedge possible.

The ask

We are raising to scale the worker fleet, harden the engine for institutional size, and pursue MM-status applications across the four venues that offer elevated rate limits. Details in §10.

Market opportunity

Perp DEX volume

Perpetual DEXes are the fastest-growing segment of crypto market structure. 2025 closed at over $2T in cumulative perpetual volume across the major venues, up from under $700B in 2024. Hyperliquid alone now routinely clears $8-12B daily — comparable to mid-tier centralized exchanges.

The growth has structural drivers, not just narrative:

- **No KYC, no custody.** Users keep their own keys.
- **Settlement transparency.** Every fill is on-chain.
- **Cross-venue programmability.** Anyone can integrate the same way an exchange operator can — there is no privileged “API user” tier.

Spread universe

There are 10 unordered venue pairs across the five integrated exchanges (Hyperliquid, Extended, Lighter, Orderly, Aster). The intersection of listed perpetual markets across pairs is in the hundreds of canonical symbols, with the top-100 by volume listed on 3+ venues simultaneously.

For each (market, venue-pair) tuple, two arbitrage modalities exist:

1. **Net-funding arb** — capture the funding-rate differential between venues. Funding settles every 1–8 hours depending on venue; rates routinely diverge by 5–50 basis points per period on liquid markets.
2. **Basis trades** — capture price-differential between venues on the same perpetual. Opens during volatility spikes, closes minutes to hours later.

The opportunity set updates continuously; Spreadr’s terminal ranks all currently-available pairs by APR-equivalent yield, with stability scores to filter for trades likely to persist.

Spreader TAM

Today the addressable users for an institutional-grade spreader are roughly four buckets:

- **Sophisticated retail / treasuries** managing 6–7-figure size, currently using manual entries or bespoke scripts.
- **Family offices + crypto funds** without dedicated trading desks but with capital and risk appetite.
- **Market-neutral hedge funds** supplementing their existing CEX arb books with DEX volume.
- **DAO treasuries** looking to put on delta-neutral yield positions on holdings.

Even a conservative penetration of these buckets at modest per-account sizes implies a multi-billion-dollar annual notional throughput.

The problem

Today's workflow is broken

To execute a cross-venue perp spread today, a user must:

1. Sign up + fund on both venues separately.
2. Generate venue-specific API credentials (different schemes per venue: agent wallets, Stark sub-keys, ed25519 keypairs, L1-derived signing keys, HMAC secrets).
3. Build or buy a bot that connects to both venues, understands both signing schemes, both order formats, both rate-limit regimes.
4. Solve cross-venue execution: monitor both books, decide when to enter, fire both legs without leaving themselves one-legged.
5. Handle disconnects, partial fills, exchange outages, funding-rate inversions — at 3 AM, automatically.
6. Manage 5+ venues' worth of operational complexity: rate limits, maintenance windows, key rotations, version bumps.

The result: a tiny minority of users (proprietary HFT shops + a handful of manual desks) capture nearly all of the spread opportunity, while the majority of would-be participants stay out.

Specific pain points

Leg risk.
Maker fills, hedge fails. Sudden directional exposure. Most retail tools have no answer; manual desks scramble.

Key management.
API keys with trade-only scopes are venue-specific, often complex (e.g. HL agent + parent), and rarely stored with KMS-class isolation.

Crash recovery.
A worker dies mid-trade. Did the hedge send? Is the position one-legged? Reconciliation requires manual review.

Rate-limit contention.
Run two spreads on the same wallet → both share the same per-IP rate bucket → both throttle. No clean isolation in DIY setups.

Latency.
Spread opportunities decay in seconds. Bots running on US east-coast home connections lose to anything Tokyo-colocated.

Observability.
DIY tooling gives raw exchange logs. Net PnL across two legs, accrued funding, delta drift — all assembled by hand.

Why this hasn't been solved

The technical surface is genuinely difficult: five venues, five signing schemes, five rate-limit regimes, four order-placement transports. Most bot vendors stop at one or two CEXes. Most DeFi tooling stops at single-venue execution. The gap is structural, not just a missing product.

The Spreadr solution

What it is

Spreadr is a SaaS execution platform that productizes everything above. A user signs in once with their wallet, links credentials for each venue they want to trade on, and configures spreads in a web terminal. The execution engine — running co-located in Tokyo — handles every leg.

System shape

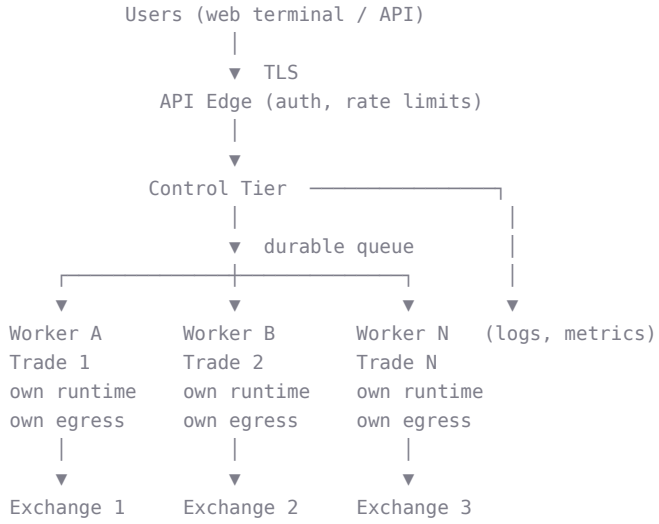

The five guarantees

01 **Isolation.** Each spread runs in its own runtime with its own network egress. Per-trade rate-limit budgets, no cross-contamination between users.

02 **Cryptographic credential isolation.** Credentials are encrypted under a user-scoped encryption context. Decryption fails if the context doesn't match — even a misconfigured access policy cannot reveal one user's keys to another.

03 **Fault tolerance.** Every fill is durably logged **before** its hedge is sent. On worker crash, a replacement reconciles the ledger against the exchange — orphan orders cancelled, unhedged fills caught up. Never double-hedge.

04 **Strict risk controls.** Per-exchange leverage caps, margin pre-flight on every order, chunked execution for large sizes with margin re-check between chunks, global + per-user kill switches.

05 **Full audit trail.** Every fill, hedge, reject, and amendment is logged with a sequence number. Users can request the audit log for compliance or 4-eyes review.

Why hosted, not self-hosted

Latency requires colocation. Security requires server-side signing. Reliability requires a 24/7 worker pool. Self-hosting any one of these is hard; self-hosting all three is infeasible for most users. Spreadr's

operating cost is amortized across all users; each user gets institutional-grade infrastructure at a fraction of the cost of building it.

Why this doesn’t exist yet

We have looked carefully and have not found a productized, institutional-grade automated cross-DEX-perp spreader on the market today. Adjacent categories exist; none cover the same surface.

Adjacent categories

Funding-rate analytics platforms (e.g. Pear Protocol’s market view, CoinGlass, Funding-Rates.io). Show opportunities but do not execute. The bot is the user’s problem.
Open-source bot frameworks (e.g. Hummingbot). Execute, but require the user to self-host, write venue integrations, manage credentials, and operate their own infrastructure. The “Spreadr is hosted” guarantee is gone.
Manual prop desks. High-touch, high-capital. The current capture mechanism for most institutional-size cross-venue arbitrage. Doesn’t scale to retail or smaller funds.
Build-your-own. Every sophisticated trader who’s tried this. They’ve each rebuilt the same venue integrations, rate-limit handlers, and reconciliation logic — usually badly, always at private cost.

Spreadr’s position

Spreadr is the first productized solution that combines all four: analytics (we surface the opportunities), execution (we run the trades), managed (we operate the infrastructure), and multi-venue (five exchanges integrated end-to-end). The categories above each cover one leg of that. Spreadr covers the whole stool.

Defensibility

We expect competition to emerge — the category is too large to stay empty. Spreadr’s defensibility:

- **Integration depth.** Each new venue is months of integration work. Spreadr has done it for five and has a plugin framework that makes the sixth marginal.
- **Operational track record.** Crash-recovery semantics and never-double-hedge guarantees take real production incidents to harden. We have lived through them.
- **Latency moat.** Tokyo colocation + per-trade isolation is operationally non-trivial. A self-hosted clone would have to replicate the entire infrastructure layer.
- **Builder relationships.** We are registered as a builder on four of the five venues. This is a fee-revenue moat that doesn’t go away.

Traction

Status

Spreadr is currently in an early-access cohort. The execution engine has been running in test against live exchanges since early 2026; the multi-user web terminal launched into limited access in Q2 2026.

Engine readiness

- ✓ Five exchanges integrated end-to-end (HL, Extended, Lighter, Orderly, Aster), each with their native signing scheme and lowest-latency transport.
- ✓ Maker-taker strategy battle-tested across 24+ exchange-pair combinations.
- ✓ Chunked execution validated up to seven-figure notional.
- ✓ Crash recovery + reconciliation tested in fault-injection drills (orphan orders, unhedged fills, WebSocket disconnects, exchange-side outages).
- ✓ Tokyo colocation operational.
- ✓ Per-user-encrypted credential vault deployed; control tier cannot decrypt by construction.

Platform readiness

- ✓ Production web terminal at app.spreadr.xyz, with real-time positions panel, spread chart, order book ladder, and audit log.
- ✓ REST + WebSocket API surface for programmatic access.
- ✓ Public documentation at docs.spreadr.xyz.
- ✓ Marketing site at spreadr.xyz.

Validation

The early-access cohort has been intentionally small while we harden the engine for institutional size. Cohort members have stress-tested the platform across all five venues, run real spreads against real capital, and provided product feedback that has informed every release of the engine.

What we want from this raise

Two things, and they are tightly coupled:

1. **Scale the worker fleet.** Today's infrastructure can comfortably handle the early-access cohort; we need to harden it for the first hundreds of concurrent paying users without sacrificing per-trade isolation.
2. **Open public access.** With the engine stress-tested and the platform polished, we are ready to transition from waitlist-controlled access to open-access.

Business model

Today: exchange builder fees

Spreadr is registered as a builder on four of the five venues (Hyperliquid, Extended, Lighter, Orderly) and earns a builder fee on each fill it routes for the user. The fee is taken from the exchange side, not the user side — it is a portion of the existing exchange fees the user would have paid anyway. Aster's builder program is in development and will be added on launch.

For the user, this is a **transparent line item** on every fill. For Spreadr, it scales 1:1 with platform volume.

Implication. Spreadr's revenue is structurally aligned with user activity: more spreads = more fees = more revenue, with no incentive to over-trade or push users into unprofitable positions.

Scale economics. Builder fees are predictable per unit of volume. Doubling the user base roughly doubles revenue, with infrastructure cost scaling sub-linearly (per-trade workers are billed by the second).

Near-term: subscription tier (planned, Q3-Q4 2026)

For high-volume users (six-figure-plus monthly notional), we will offer an optional subscription tier that includes priority worker allocation, elevated rate-limit budgets, white-glove integration support, and custom dashboards. Pricing will be set after additional cohort feedback.

Long-term: API-product tier

The Spreadr REST + WebSocket API is already public-grade. We expect a non-trivial fraction of long-term volume to come from third-party integrators (other front-ends, structured-product wrappers, fund back-ends). API-only access will be priced as a separate tier when the integrator pipeline warrants it.

What we will not do

- **Charge users at the wallet level.** No Spreadr-side fee deducted from user funds. Builder fees are exchange-side, period.
- **Take custody.** Spreadr never holds user funds; trades execute against the user's own exchange accounts.
- **Front-run users.** Spreadr does not trade for its own book. Period.

Team

About hej

Spreadr is built by **hej**, a crypto-native engineering team based in Tokyo. The team has prior experience across DEX integration, market microstructure, low-latency systems, and trading infrastructure at both crypto-native and traditional finance institutions.

Anonymity, on purpose

Individual team members are intentionally pseudonymous on public documentation. This is standard practice for crypto-native teams whose on-chain footprint would otherwise become a target surface — every team member has live trading positions, and de-anonymization creates unnecessary operational risk.

Individual founder and engineering bios are available under NDA for serious investor diligence. Reach out via the contact channels on spreadr.xyz to arrange.

Composition

Founder.	Crypto-native trader with cross-venue arbitrage execution background; built and ran prior bots at production scale before starting Spreadr.
Engineering.	Distributed-systems engineering experience spanning low-latency execution, fault-tolerant state machines, and AWS-scale infrastructure operations.
Operations.	Trading operations experience including on-exchange relationships, market-maker programs, and builder integrations.

Advisors + ecosystem

Spreadr operates inside a network of crypto trading firms and DEX integrators that has supported the integration depth and operational hardening required for the engine. Specific advisors and ecosystem partners are not listed publicly to preserve their own discretion; references available under NDA.

Roadmap + ask

Roadmap (next 12 months)

Q3 2026	Open-access launch. Subscription tier for high-volume users. Public-access dashboard with anonymized cross-venue spread analytics.
Q4 2026	Sixth + seventh venue integration (priority: any venue clearing $5B+ daily). MM-status applications on the four eligible exchanges.
Q1 2027	Cross-asset hedging (perp ↔ futures basis, where venues offer both). Public API-product tier launch.
Q2 2027	Custom strategy SDK for power users (run user-defined strategies inside Spreadr's hosted execution + risk infrastructure).

Ask

We are raising to fund the open-access launch and the next four venue integrations. Specifically:

40%	Engineering — accelerate venue integrations, MM-status hardening, cross-asset hedging engine.
25%	Infrastructure — scale the worker fleet for open-access launch, improve observability + auto-scaling.
20%	Operations — market-maker applications, exchange relationships, compliance + risk-management buildout.
15%	Growth — early-access cohort expansion, content + partnerships, integrator relationships.

Why now

The perp DEX market has matured into a serious institutional opportunity, but the execution tooling for cross-venue capture has not followed. Spreadr is the productized solution to that gap, with the engine already production-tested and the platform already live. The window to define the category is now.

For investor inquiries: see spreadr.xyz · for technical diligence: docs.spreadr.xyz