GREEN SAND STRATEGIES

White Paper

Engine, Product, and Roadmap

EDITION

Public

DATE

August 12, 2026

STATUS

Not investment advice, not an offer to sell securities, not a prospectus

READ ONLINE

greensandstrategies.com/whitepaper/

A pre-registered trading research program, and the public product built on top of it.

> This document describes a pre-revenue, pre-entity venture. Nothing here is investment advice, an offer to sell securities, or a guarantee of future results. Every signal surface we ship carries an explicit not-investment-advice disclaimer, and the flagship product is positioned as entertainment and education.

1 Executive summary

Green Sand Strategies runs two things under one name. A private, pre-registered trading research program builds and tests a systematic strategy on the founder's own hardware. A public product turns what that research produces into something people and software agents can use.

The research engine is a multi-timeframe trend-alignment breakout system. It runs privately, stays off the cloud, and never places a trade on its own. What it produces is a stream of timestamped evidence: every decision it makes, including every time it decides not to trade, gets written to a permanent, append-only record before anyone knows how that decision turns out. That record is the company's real asset. The public product is how customers get access to what the record produces.

As of this writing, the core validation gate has cleared its first pre-registered milestone. The live signal daemon has been logging real decisions since August 8, 2026. The public site has been live since August 9, 2026, with working accounts and paper trading already in place. The business side — subscriptions, dataset sales, an agent API — is still ahead. That is the roadmap this document lays out.

2 The thesis

Green sand is a foundry term: a mixture of sand, clay, and water, packed around a pattern to cast molten metal into a finished part. A foundry's value lives in two places at once. There is the pattern shop and the metallurgy — the proprietary process that turns raw material into something with the right strength and finish. And there is the catalog of finished castings that customers actually buy, rent tooling against, or build machines around. A customer needs the part to be reliably good. They do not need the furnace recipe to use it.

That split is the whole company. Two convictions sit underneath it.

Evidence has to be pre-registered to mean anything.

> A backtest optimized after the fact proves that the researcher had enough parameters to fit the past, nothing more. Our engine commits to a hypothesis, a scoring rule, and a falsification condition before seeing forward data, then logs every live decision — trade or no-trade — to an append-only file it cannot go back and edit. Several parameter changes that looked obviously correct have already been tested this way and rejected; the strategy configuration carries that falsified-fix record inline, on

purpose, so nobody quietly re-adds a gate that already failed. That discipline is slower than most retail systems build for. It is also the entire point.

The research process and the customer-facing product work best as two separate businesses.

The signal engine loses its value the moment its internals go public: the measures it uses, their parameters, position sizing, and the full scenario logic stay on private hardware, permanently. What reaches the public is a derived, sanitized signal contract — mapped states like `signal`, `neutral`, `vetoed`, and `blackout` — with the underlying risk math stripped out. Customers buy or rent access to what that process produces. The research stays in the foundry; the castings go out the door.

Run well, this becomes a flywheel. Rigorous, falsifiable research produces evidence worth paying for. Revenue from that evidence funds more research and a longer validation history. A longer, cleaner validation history makes the next product easier to sell on demonstrated trust rather than promises. Each turn of the cycle makes the catalog more valuable without ever opening the foundry floor.

3 The engine — the foundry floor

Everything in this section is the part of the business that never ships directly to a customer. It is worth describing anyway, for the same reason a foundry describes its metallurgy to a customer without handing over the furnace: the quality of the process is what makes the catalog worth trusting.

Strategy

A trend-alignment breakout system built from three independently swappable modules: entry scoring, participation confirmation, and exit scoring. Which measures fill those slots, and the parameters behind them, are not published.

Regime gate

A higher-timeframe veto sits in front of every entry, with telemetry so the gate's own performance gets measured over time, the same as everything else.

Signal daemon — live since August 8, 2026

Scores the watchlist, applies the regime gate, checks earnings blackouts, and logs every decision — including no-trades — to an append-only record. It places no orders and cannot reach an order endpoint. Any future execution layer will be a separate, separately-armed service.

Validation program

Monte-Carlo low-timeframe backtests, circular-block and cluster-robust confidence intervals, a reproducible go/no-go bootstrap gate, and an explicit overnight gap-through-stop risk measurement. This is instrumentation built specifically to make it hard to fool yourself.

A second research line

A complementary pre-open research line has independently cleared its own pre-registered validation gate and now runs on an automated weekday schedule. It is a standalone toolset, still being built out,

and it has not been merged into the primary signal pipeline. Even so, it is a second data point that the validation methodology generalizes beyond the first system.

> *Every claim in this document about the engine can be verified against an append-only log with a timestamp older than the claim itself.*

The method note at greensandstrategies.com/method/ goes further, including a headline finding we withdrew after discovering it was an artifact of our own label definition.

4 The product — the storefront and catalog

If the engine is the foundry floor, the app is the storefront and the catalog combined — where what the research produces reaches a customer, human or software, ready to use. There is an analogue-style paper-trading app for people at greensandstrategies.com and a parallel API for agents at api.greensandstrategies.net, both live since August 9, 2026.

Architecture

A single-operator-friendly, all-Cloudflare stack: Workers for the API and auth, Pages for the frontend, D1 for accounts and the paper-trading ledger, Workers KV for cached signal snapshots, WAF rate limits for abuse protection, and cookieless analytics. R2 is reserved for dataset storage because its egress is free; the first published dataset currently ships as a static asset while R2 billing is still pending. The private signal daemon pushes sanitized snapshots into this stack over an authenticated ingest endpoint. Engine internals never travel as code or as data.

Cost discipline

A hard $50/month infrastructure ceiling, enforced by free-tier serving caps, zero-egress storage, no autoscaling dependency, no payment method on file with any provider that could run away on cost, and billing alerts at $25 and $40. The business can only spend within limits it chose in advance.

Shipped and live today

- Public signal showcase — the sanitized, delayed signal feed, live to anyone, no account required
- Passkey accounts — no password, no email
- Event-sourced paper-trading ledger — a full trade history model, no real money in the loop

Not built yet

- Stripe monetization: real-time signals, dataset sales, metered agent API keys
- MCP server and agent marketplace

5 Business model

Free tier

Delayed signals — currently 15 minutes — cookieless, no account required.

Paid tier (ahead)

Real-time signals and reference levels. The delay itself is the monetization lever: free stays genuinely useful, real-time is the upgrade.

Dataset sales

Derived analytics only, never raw or reconstructable market data. That boundary is a deliberate licensing choice, built into the product from the start.

Agent API (ahead)

Metered keys, aimed at the MCP and agent-tooling ecosystem specifically.

No pricing has been set, and no payment infrastructure exists in code. This is honestly the least-built part of the company today.

6 Traction, plainly

The previous pipeline was retired on August 8, 2026, the same day the current scenario engine and its signal daemon went live and began logging forward evidence. A day later, on August 9, 2026, greensandstrategies.com and api.greensandstrategies.net launched together, with passkey accounts and the event-sourced paper-trading ledger shipping in the same release. Since then the pre-registered validation program has kept running, and the second research line cleared its own validation gate independently and moved to a scheduled, automated live run.

That is the complete list. There is no revenue yet, no paying customer, and no payment infrastructure built. Saying so plainly here is what makes the rest of this document worth trusting.

7 Roadmap

Near term — monetization

Ship billing, the real-time paid tier, and the first dataset-shop listings. The gate for turning this on: the free-tier signal contract needs enough forward history that a paying customer is not the first person testing it.

Following — agent marketplace

Ship the MCP server and the metered agent API. This is the bet that a meaningful share of future demand for market signals will come from other software rather than people, and that being agent-

native from the start is a real advantage over incumbents retrofitting an API onto a human-first product.

Ongoing — engine side

Keep the validation program running as the core of the moat: every strategy change stays pre-registered, every forward log stays append-only, and falsified ideas stay documented rather than quietly retried. Expand the dataset library and bring the second research line up to the same standard as the primary engine, so it can eventually feed the product too.

Commitments that hold regardless of phase

The private engine stays off the cloud, permanently. Any future execution layer stays separate from the research engine and gets armed on its own, deliberately. Every public signal surface keeps its entertainment-and-education framing as the product grows.
