The Meridian Position

What is held, what it proves, what it earns, and the route from a platform to a fund

Meridian Deal Partners · 1 November 2026

This document states the position completely, including the parts that did not work. Three of the four research signals we carried died under proper controls; the allocator outreach lane has been running at roughly a quarter of one per cent; one of the seven sponsors we went looking for publishes no leadership page at all. Those facts are here for the same reason the successes are — a position you cannot check is not a position, it is a pitch.

1 The position in one page

Meridian holds 5,058,265,273 records drawn from public registers, of which 1,864,734,940 are joined to a single identity spine. On top of that estate sits a research programme that has been tested adversarially, an origination layer over physical energy infrastructure, a book of 663 named decision-makers at institutions above a billion dollars, and — as of this week — a working path to invoice and collect.

The commercial question is which of three engines to run first, and one external number settles it. BlackRock paid 13 times revenue for Preqin — $3.2bn against roughly $240m of recurring revenue. In this business a dollar of recurring revenue is therefore worth about thirteen dollars of enterprise value, which means the engine that pays the bills and the engine that builds the valuation are the same engine. The sequence follows from that and not from preference.

The second thing that settles the sequence is internal. Of four candidate trading signals, three lost between 75 and 82 per cent of their effect once characteristic-matched. The survivor is a crash-risk signal rather than a return forecast. Read as a strategist rather than a researcher, that says the engine is materially better at origination than at prediction — so leading with a systematic fund would put the weakest work at the front and invite the diligence that breaks it.

Table 1. The position at a glance. Each line is measured and each is expanded below.

The position	Measured	Status
Estate, resolved to identity	1,864,734,940	built and maintained
Replication gate	+8.16% vs ~7%	passed
Signals surviving controls	1 of 4	a risk signal, not a forecast
Firm capacity, verified	18 of 372	the scarce commercial answer
Named $1bn+ decision-makers	663	11 approached
Sponsor contacts discovered	57	this week, 6 of 7 sponsors
Billing path	client to cash	built, tested, scheduled

2 The estate: what is actually held

The estate is 5,058,265,273 records. The number that matters commercially is smaller: 1,864,734,940 are joined to a single identity spine, and 3,193,530,333 are held adjacent for reference and never mixed into an analytical claim. Only the resolved half can be sold as an answer, and keeping that boundary visible is what makes the rest credible.

It spans 28 lake corpora and 33 local stores across 263 tables, maintained by 299 scheduled jobs. The substantive holdings include 590,156 Form D private-placement filings, 17,566,355 reported 13F holdings, 35,878 manager anomaly scores, 312,377 dependency-graph edges, and a relationship layer of 17,090 organisations with 25,417 contacts.

What makes it an asset is not the volume. It is that four decades of disclosure law pushed an enormous quantity of economic detail into the public record, each register built by a different institution for its own purpose and none designed to be read against another. The opacity is not secrecy; it is a filing problem, and a filing problem is an engineering problem. The estate is the answer to that problem, and its replacement cost — a multi-year engineering programme — is the honest valuation anchor at this stage.

3 What the research licenses us to claim

A research programme that reports only its successes cannot be checked, so the pipeline was first pointed at a result somebody else had already published. Brav, Jiang, Partnoy and Thomas reported roughly a seven per cent abnormal return around an initial Schedule 13D, and no subsequent reversal.

Table 2. The research record, successes and failures together. Recovering both the magnitude and the absence of reversal, from a different period and on independently assembled filings, is stronger evidence than matching a headline.

Test	Result	Reading
Replication, (–20,+20)	+8.16% n=266 t=2.63	published benchmark ≈ +7%
Reversal, (+1,+60)	+0.21% t=0.09	no reversal — their central claim
Signals before controls	4 of 4 strong	t between 3.1 and 11.9
Signals after matching	1 of 4	three lost 75–82% of effect
The survivor	–15.15% t = –6.03	a short; tail-driven, skew –2.78

Note: the survivor's median name is positive. The effect is a minority of severe collapses, so it reads as crash risk and must never be presented as a return forecast.

Two honest limits travel with this. The lab has never run genuinely out-of-sample — until a forward window is scored on data the models never saw, what exists is a strong in-sample result with hostile controls, which is a different and weaker claim. And matured predictions are not yet scored with their misses included. Both are gates to the fund, and both are stated here because an investor who discovers a missing gate assumes there are none.

4 The origination layer, which is the sharpest thing we hold

If the engine is better at origination than prediction, the question becomes what it originates that nobody else can. The answer is physical energy infrastructure, in the sector absorbing more capital than any other on earth.

Of 372 datacentre developer plays tracked against the public register, 18 carry firm capacity rather than a queue position — 4.8 per cent. In aggregate that is 196,055 MW announced against 32,047 MW confirmed. The gap between announced and confirmed is the question that market turns on, and we hold both sides of it. Beneath it sit 46,434 federal power and grid awards resolving to 4,231 named suppliers with $12.1bn obligated, and 11,434 generators joined from physical asset through owner and legal-entity identifier to the listed holders exposed to them.

Resolving those plays to their actual counterparties was the work of this week, and it pruned hard. The plays are project vehicles, so the sponsor behind each one is the counterparty rather than the entity on the register. 18,195 MW sit behind listed issuers — Fermi, Vistra, NRG — and are excluded, because a listed issuer raises in the public markets, which is a different and more regulated conversation. What remains is 6 addressable sponsors carrying 10,156 MW.

Table 3. The addressable sponsors after exclusions, with the 57 contacts discovered this week. Five of six now have a named chief executive or founder.

Sponsor	Firm MW	Contacts found	Note
Knighthead / Homer City	4,396	21	CEO and full partnership
CPV	1,445	5	CEO named
Invenergy	1,288	10	founder named
EmberClear / Lincoln Land	1,277	0	publishes no leadership page
Copia Power	1,250	11	CEO and CIO named
TerraPower	500	10	CEO named

Note: contacts are candidates from each firm's own published leadership page and are held for review. None has been contacted, and a generic inbox is never an acceptable substitute for a named person.

5 Distribution: who we can actually reach

Across all lanes 23,609 approaches have been sent and 424 have drawn a reply — 1.80 per cent. The allocator lane specifically has sent 7,563 and drawn 20 — 0.26 per cent, roughly a seventh of the book's own rate.

That gap is diagnostic and it is not about the copy. Those sends went to Form ADV registrants above a billion dollars, which turns out to mean retail advisory aggregators and robo-advisers, and to 13F cover-page signatories, which are 67 per cent compliance and operations officers because a 13F is signed by the chief compliance officer. A good message to the wrong title produces exactly that number, and no rewrite moves it.

The replacement is a two-tier book. 663 named decision-makers at institutions above a billion dollars, with compliance and operations stripped out and every row carrying a title, a direct line and the filing it came from — segmented into 29 allocators (pensions, foundations and insurers, including the Gates Foundation Trust, Ohio PERS, the New Jersey Common Pension Fund and Fjärde AP-Fonden), 6 quantitative shops, and 628 platforms. This tier is worked by hand and cannot be bulk-mailed: allocators are a referral graph, and volume is precisely what produced the quarter of a per cent.

6 The three engines

There are exactly three ways to convert this estate into money. They differ less in size than in how controllable they are, which is the property that should decide sequencing.

Table 4. The three engines, ordered by time to first cash.

Engine	To cash	Ticket	Gated on
Subscription intelligence	weeks	$25–75k / yr	a billing path and one priced product
Advisory / capital intro	months	5–6 figures / deal	an agent of record
The fund	2–4 years	2% + carry	out-of-sample proof; a record substitute

Engine one. Pricing anchors are institutional — Enverus and Wood Mackenzie in energy research, Preqin in private markets. Ten clients at $50k is half a million of recurring revenue and roughly $6.5m of enterprise value at the Preqin basis; forty at $60k is $2.4m and roughly $31m. The leverage is in price rather than volume — moving the ticket from $50k to $75k is a fifty per cent revenue increase at no extra sales effort, and scarcity (4.8 per cent of tracked developers clear the firm-capacity test) is exactly what supports the higher number. Underpricing is the easiest avoidable mistake available to us.

Engine two. We hold both sides of a transaction: 6 sponsors with 10,156 MW of confirmed capacity on one side, 663 named institutional decision-makers on the other. It is the highest-margin position in the business and the least controllable — one transaction can outweigh a year of subscriptions and can equally slip two quarters for reasons unrelated to us. It is upside, never the base, and it requires an agent of record before it can transact at all.

Engine three. Aldaron raised approximately $885m for a debut fund in October 2026 on an engineering capability and a narrow mandate, ahead of any investing record. That is proof the shape works. Their substitute for a record was Palantir and Carlyle pedigree; ours has to be the estate, the replication gate, and a sourcing register that includes its dead ends. That difference is real and the plan has to close it rather than ignore it.

7 Comparables, read properly

A comparable is only worth citing if it identifies the metric that drives the value. The commonest failure in a document like this is borrowing a headline multiple and applying it to a business measured on something else entirely. Three destinations exist and all three are priced differently.

Table 5. The three exits and the metric each one actually turns on.

Destination	Precedent	Paid	Priced on
Sell the business	Preqin → BlackRock	$3.2bn	~13× recurring revenue
Licence the method	RAFI → TMX Group	$490m	$147bn tracking it
Capability to carry	Aldaron, debut fund	$885m	legibility, not prior returns

Note: Preqin is a revenue multiple, RAFI is an assets-tracking business, Aldaron is a commitments business. Only the third is reachable without first manufacturing years of recurring revenue — which is why Aldaron chose it over selling engineering by the hour.

Read the table one way and it is a menu. Read it the other way and it is an instruction: the comparable tells you which metric to manufacture. Preqin says build recurring revenue. RAFI says build a licensable methodology and find a benchmark partner. Aldaron says make the capability legible and the mandate narrow.

8 From platform to fund: the sequence, and why it is not a compromise

Build the subscription engine first. The argument is not that it is a modest start. It is that it produces four things simultaneously, where no other engine produces more than one.

One. Cash, in weeks rather than years.

Two. Enterprise value, at roughly thirteen times whatever it earns.

Three. Paying institutional clients — which is itself the track-record substitute. Six infrastructure funds paying for this is a stronger opening with an allocator than any backtest, and it is the specific thing Aldaron had and we do not.

Four. The distribution a fund raise requires, built in advance and paid for by someone else.

Going straight at the fund inverts all four: a long runway with no cash, no valuation anchor, and a diligence conversation we are not equipped to win until the out-of-sample work lands. So the fund is funded out of the subscription engine and the research gates proceed on their own clock. If they land, we raise with revenue, clients and proof. If they do not, we still own a business worth a multiple of its revenue. That asymmetry is the entire design, and it is why the sequence is a strategy rather than a concession.

9 What is built and running

The operating system is 299 scheduled jobs over 263 tables. Three pieces changed this week and they are the ones that matter commercially.

The billing path now exists end to end: client, engagement, invoice, settlement and a recurring-revenue line, with sequential invoice numbering and an idempotent cycle. Two rules are built into the code rather than left to discipline — nothing in it charges anyone, because payment links stay behind an existing human gate, and settlement reconciliation is read-only against Stripe, so an invoice being marked paid is the consequence of a settlement rather than the cause of one. The annual intelligence tiers are priced in the catalogue at $25k, $50k and $75k alongside the existing monthly retainers.

Outbound was repaired. The follow-up lane had been silently dead for four weeks: a single declined recipient sat at the head of a one-row queue and was re-selected every run, holding 6,798 due follow-ups behind it while the log reported an empty queue. It is sending again — 176 approaches went out today. And

contact discovery, which is the constraint on every route at once, went from nothing to 57 named sponsor contacts.

That last one is worth a sentence on method, because the first attempt failed instructively. A regular expression over page text returned 220 candidates of which roughly 97 per cent were organisations and places — New Mexico, United States Navy, Georgia Institute — the identical failure that had earlier put entries like Managing Member into the advisory book. Names are found structurally or not at all.

10 The gates, and what would falsify this plan

Stated as conditions that can fail, because a plan with no failure conditions is a wish.

Table 6. The five gates. The third and fourth decide whether the platform business is real; the first two decide whether the fund is reachable.

Gate	Why it matters
Out-of-sample scoring	never run; until it is, the research claim is in-sample
Live prediction resolution	misses must be scored too, or the learning loop never closes
Contact discovery at scale	the live bottleneck; blocks all three engines at once
First three paying clients	proves the product is a product and not a capability
An agent of record	without it the highest-margin engine cannot transact

What would falsify the plan: if twenty to thirty qualified conversations produce no paying client, the product is not a product and the pricing or the packaging is wrong rather than the market. If the out-of-sample window comes back flat, the fund route closes and the business is a research and origination platform, which is a smaller but entirely real outcome. Either way the failure is visible early and cheaply, which is the point of running the engines in this order.

11 The plan

The next ninety days are about the first three paying clients rather than about scale. On a ninety-day enterprise research sale and a ten to fifteen per cent conversion from a qualified named conversation, three clients implies twenty to thirty real conversations. At $50k that is $150k of recurring revenue booked, which on the thirteen-times basis is roughly $2m of enterprise value created in a single quarter.

Through month twelve, two to three new clients a month is the normal rate for a product at this price with a single seller — $1.2m to $1.8m of recurring revenue, or roughly $16m to $23m of value. Advisory covers the gap irregularly and should not be planned around. The research gates run in parallel on their own clock throughout, because they are slow and because every serious allocator tests them first.

Those are scenarios rather than forecasts, and each assumption sits next to the number it produces so it can be argued with directly. The honest summary is narrower than the document: a billing path, one priced product, and contact discovery. The first is built. The third moved this week. The second is the next thing to make.

Comparables: BlackRock investor presentation on the Preqin acquisition (13× 2024E revenue) and S&P Global Ratings; TMX Group agreement to acquire RAFI Indices; Private Equity Wire on Aldaron's debut fund. Research figures from figures.json; operating figures measured from the live stores on 1 November 2026. Sponsor contacts are candidates from each firm's own published leadership pages, held for review and not yet approached.