




The investment ecosystem for operating real estate

One record, many readers. It lives with the property forever.
forever.

Phil Bernardo, Founder · Pre-Seed · September 2026



THE PROBLEM

The same asset gets modeled six times by six parties who never reconcile

WHO	WHAT THEY DO
Broker	builds the marketing proforma
Sponsor	rebuilds it to bid
Lender	rebuilds it to size debt
Insurance	rebuilds it to quote
IC	rebuilds it to approve
Asset management	rebuilds it to report

Six versions, no lineage

When the deal underperforms, nobody can attribute the miss to the assumption that caused it, because that assumption no longer exists in a form anyone can query.

The model is abandoned at close, the moment it starts to be worth something.

40:1

Deals underwritten for every one closed. The rebuild happens forty times before a single asset trades.

The real incumbent is an abandoned spreadsheet.

WHY IT IS STILL UNSOLVED

These assets price by the night, not by the lease



HOTELS



SHORT-TERM RENTAL



SELF STORAGE



SENIOR HOUSING

ASSET CLASS	REVENUE UNIT	CADENCE	WHAT A LEASE ENGINE CANNOT REPRESENT
Hotels	RevPAR per available room	Nightly	Departmental P&L, franchise and management fees, PIP
Short-term rental	Revenue per available night	Nightly, by channel	Channel mix, take rates, turnover cost, no tenancy
Self storage	RevPAF per square foot	Monthly, no term	Unit mix, existing-customer rate increases, rate spread
Senior housing	RevPAU per unit plus care	Monthly plus acuity	Licensed care levels, staffing-driven expense, census

Six layers around one decision record




WHY IT IS ONE SYSTEM

Each turn sharpens the next

Signaling finds the asset, brokers bring the book, the Studio prices it, the marketplace funds it, asset management holds the asset to that price, and the data layer makes the next underwriting sharper.

Nothing leaves the wheel. Every turn adds approved assumptions and realised outcomes to the same record, so the next turn costs less than the last.

**The value is the absence of re-keying:
no layer retypes another layer's
numbers.**



Underwriting is a data-entry problem wearing an analysis costume

01 Read

The language model lifts the rent roll, operating statement and STAR data from the OM and tags every value with its source page.

02 Compute

The Python engine builds the USALI departmental proforma, sizes debt on DSCR, debt yield and LTV, and returns a ten-year cash flow.

03 Check

The coherence engine flags assumptions that disagree with each other or with the market panel. Every override needs a written reason.

04 Hold

The run persists. The broker package, credit memo, IC deck and variance report are views of that run, never copies.

EVIDENCE OF BUILD

Engine roughly 90% complete, running against live offering memoranda today.

147 brands across 12 franchisors in the fee repository, on 2026 FDD vintages.

Same inputs, same outputs, every run. A test keeps the AI service out of the compute layer.

THE DECISION LAYER

The agent decides what to ask. The engine decides what is true.

Every proptech deck says AI-powered. PATL moves the language model further from the number.

MCP TOOLS ON THE SERVER

run_deal

update_assumptions

get_deal_diagnostics

get_brand_fees

get_market_snapshot

list_deals

write_review_memo

Fifteen tools in total. None of them can return a cash flow the engine did not compute.

The constraint

MCP tools call the compute layer and never the language model layer. A test enforces it.

An agent that asks PATL for a number gets the same number every time. That property is what lets parties who do not trust each other share one model.

AGENT AGNOSTIC

ChatGPT

Claude

Perplexity

Gemini

PATL is a tool, not a chat product. Any MCP client calls the same server and gets the same numbers, so a customer keeps whichever assistant they already use.

Flow arrives before the book goes wide

Brokers give early looks to buyers who answer fast and honestly. With extraction and screening already mechanical, a credible answer takes hours.

What the broker gets

- A proforma builder that works at the listing pitch, before an OM exists.
- A live model sent to the buyer list in place of an emailed PDF.
- Pass reasons recorded against the run and returned, which earns the next early look.

Why the buyer trusts it

- The broker cannot edit the math. Broker-asserted inputs render apart from engine figures on every export.
- Every override carries a written justification that stays on the record.
- Neutrality is a property of the engine, enforced by test.

The cheapest seat in the product is the channel for the most expensive one.

Every counterparty quotes the same numbers

THE CREDIT COMMITTEE QUESTION

What changed since version one, who changed it, and why?

A workbook cannot answer it. PATL can, because an override without a written justification cannot be saved.

REQUEST

Debt and equity requests assemble from the approved run, and insurance quotes off the same property record.

QUOTES

Term sheets stay attached to the run they priced. Change an assumption and every quote it invalidates is flagged, with a dated log of who saw which version.

SCOPE

PATL routes the request and keeps the record. It does not hold the paper or write the policy.

The asset is held to the plan it was bought on

ILLUSTRATIVE

Bought at 1.41x DSCR on \$91.40 RevPAR and a stabilized \$1.84M NOI.

Twenty-six months later, is the asset on its thesis?

OPENING BUDGET	The acquisition underwriting becomes the opening budget. There is no second workbook.
VARIANCE	Monthly actuals post against that plan, with variance reported line by line.
FORECAST	Forecasts re-run on the same engine, so basis and remaining-year view stay comparable.
PORTFOLIO	Owner, board and covenant reporting roll up from the same runs across every asset held.

Actuals come from the systems that already hold them. Only the platform that underwrote the deal can report against it.

Decisions and outcomes, linked at the moment of decision

Data Intelligence

FOUNDATION BUILT

- Every run retained, won or passed, with its approved assumptions attached.
- Comparables built from reviewed inputs, queryable by market, brand, keys, vintage and basis.
- Improves each time an analyst approves a line, at no added cost.

Distress Signaling

VISION

- Coverage, basis, maturity and operating drift watched together across the tracked set.
- Thresholds set from underwriting, so a signal means something specific about one asset.
- A dated signal trail, so the earliest observation is recoverable when the asset trades.

**ALREADY
BUILT**

A license-class flag on every series row separates public federal data from licensed passthrough. This layer is excluded from every projection in this deck.

WHY NOW

Three things became true at once

1

Extraction finally works

Reading a 60-page offering memorandum with a departmental operating statement and a STAR report was not reliable two years ago. It is now, which removes the manual step that made this category uneconomic to serve.

2

Agents arrived with nothing to call

Acquisitions teams are being handed agents with no deterministic tool for operating assets to point them at. That gap stays open only until someone builds the engine.

3

The standards are already written

USALI and its analogs in storage and senior housing make the departmental structure global. The engine is specified once per asset class, not once per market.

The middle one is the timing argument. The other two are the feasibility argument.

WHY ME

I built the model. I found out it was wrong. I signed on it anyway.

01 · BUILT IT

Operations and revenue management

Marriott, Universal Orlando, Disney Parks. Ten years pricing the night and running the departmental P&L.

In the product: RevPAR modeled the way an operator manages it.

02 · FOUND OUT IT WAS WRONG

Asset management and capital markets

Xenia Hotels & Resorts, Hillpointe, Hyatt. Held assets against assumptions nobody could still explain.

In the product: the override justification field, built because I needed it.

03 · SIGNED ON IT ANYWAY

Principal and customer zero

Growth strategy director at a major hotel brand. PivotPt Capital, my fund, is the first user.

In the product: the commercial gate excludes my own network.

MARRIOTT · UNIVERSAL ORLANDO · DISNEY PARKS · XENIA · HILLPOINTE · HYATT

FULL-TIME AT FIRST CLOSE

BUSINESS MODEL

One engine, one rate card, one motion

Every tier runs the same engine. What changes is how many assets, how many readers, and how much of the record persists.

PROFORMA

\$99

Per seat, per month. Brokers, appraisers and credit analysts. The channel tier.

OPERATOR

\$499

Per month plus \$75 per asset. Owner-operators holding a portfolio.

FUND

\$90K

Base, then AUM-banded. Lender and IC surface included.

BLENDED ACV

\$45,886


Across the expected customer mix. Universe-weighted across all 42,798 owner entities it is \$11,866.

Also on the rate card: investor-host seats at \$49 a month, the asset management module at \$125 per asset a month, and broker listing packages at \$2,500 each.

Four hotels cost an owner \$9,138 a year, roughly \$150 per hotel per month. Argus Enterprise runs about \$10,000 per seat per year before anything is modeled.

DISTRIBUTION

Brokers hand the tool to buyers every time they send out a deal



WHERE WE ARE

Pre-revenue. The engine runs against live offering memoranda today, and the first paid accounts are what this round buys.

THE COMMERCIAL GATE

25 paying accounts, 15 of them outside my network. Network deals prove the product works; the other fifteen prove it sells.

WHY IT COMPOUNDS


Each layer makes the record more valuable to every party already on it, because the alternative to a shared record is a re-key.

No paid acquisition step anywhere on this slide. The broker's distribution is the channel.

MARKET

\$1.07B of software. We enter through the \$343M hotel wedge.

Bottom-up from sourced asset counts, resolved to buying entities rather than buildings: 62,000 US hotels resolve to 17,416 owner entities.



RISK

The five objections, answered

OBJECTION	ANSWER
Six layers is too much surface for a pre-seed	The layers are views over one record that already exists, and the four classes share one engine. The marketplace and signaling layers are unfunded in this round.
A foundation model will do this natively	A model can write the arithmetic. It cannot be the system of record, because a record has to return the same answer twice.
Argus and Dealpath will move down	Both are lease-based, so reaching these classes is a rewrite. Dealpath shipping MCP validates the category, and that same gap makes incumbents likelier buyers than builders.
Brokers will not adopt a neutral tool	Neutrality is what makes the buyer keep the file, and buyers discarding models is the broker's problem. The \$99 seat sells against that.
Solo founder, no commercial hire	I move full-time at first close. Twelve founder-led months, and the gate is 25 paying accounts, 15 outside my network, before any hire.

Every answer above is architecture.

THE ASK

Raising \$1.25M pre-seed

Post-money SAFE. A \$750K minimum close funds twelve founder-led months and the commercial gate without a commercial hire.

FY1 ARR

\$840K

FY3 ARR

\$8.09M

USE OF FUNDS	AMOUNT
Go-to-market and customer acquisition	\$440K
Data licensing, STR and Kalibri	\$250K
Founder salary	\$180K
Operations, tools and travel	\$220K
Engineering and contract build	\$100K
Legal, entity and insurance	\$60K
Total	\$1.25M

SALES LEADER

OPERATIONS LEADER

The \$660K across go-to-market and operations carries both hires, made after the commercial gate clears.

Distribution is funded four to one.

\$440K of reach against \$100K of code. The engine runs on live offering memoranda today, so this round buys distribution rather than build.