

A FIELD MANUAL FOR FOUNDERS AND CEOS

The Platform Value Playbook

What your real estate platform is worth, and the system that raises it.

Chirag Hathiramani | Real Estate Consulting

Creator of The Platform Value System · Former Chief Investment Officer · Q3 2026 · First Edition

Contents

	How to Use This Playbook	3
	Executive Summary	4
1	The Market That Prices You Now	6
2	The Natural Experiment	11
3	The Founder Discount	16
4	Durability and the Multiple	20
5	The Ten Dimensions Buyers Underwrite	26
6	The Seven Steps	31
7	Stake-Readiness: The Buyer's Checklist, Run Early	46
8	The Self-Diagnostic	50
	Sources and Methodology	53

This playbook is free and complete. Nothing in it is gated, and nothing essential is held back. You can run everything in these pages with your own team and never contact me. That is by design.

Your track record built the funds. This is about the firm.

If you run a real estate platform with roughly \$500 million to \$5 billion in assets, somewhere between ten and seventy-five people, and a sector thesis you would defend in any room, this playbook was written for you.

It makes one argument, and then it gives you the working method behind it. The argument: your platform, the firm itself, now has a market price, set by a class of institutional buyer that did not exist at scale a decade ago. The method: seven steps, in a specific order, that raise that price. The order matters more than any single step.

Three ways to read it. The CEO with four minutes should read the Executive Summary and the exhibit on page 5, then hand the rest to the CFO. The CFO or COO should read Sections 3 through 5, which show how a buyer constructs the number and where it leaks. The founder thinking about the next decade, a stake, a succession, employee equity that means something, or deliberate independence, should read the whole thing, and Section 7 twice.

Two commitments about what follows. First, every market figure carries a source and an as-of date; the figures will age, and the mechanisms they illustrate are the point. Second, this playbook acknowledges what the system cannot do. It multiplies a franchise. It does not resurrect one. If your sector thesis is broken, no operating layer will save the next fund, and an honest advisor will tell you so in the first meeting.

A note on vocabulary. Throughout, I say platform because that is what you built and how the buyers themselves talk. When precision requires it, the entity being priced is the management company, the firm that earns the fees, and its value is measured in fee-related earnings, FRE. That technical layer appears where methodology demands it and nowhere else.

I spent twenty years on the principal side, roughly three billion dollars of transactions, including the CIO seat at a fund manager where I raised the capital and built the platform. The frameworks here are documented at book length in *The Platform CEO* (theplatformceo.com). The playbook stands alone.

EXECUTIVE SUMMARY

Your platform has a market price now.

Over the last several years, a class of institutional investor put a public price on real estate platforms like yours. Almanac, Bonaccord, Kudu, Cantilever and their peers have bought minority stakes in dozens of mid-market managers, and the arithmetic they use is knowable: a platform is worth a multiple of its durable fee-related earnings. In today's mid-market, that multiple runs roughly eight to twelve times.¹

Read that arithmetic once more. It means every dollar of durable platform margin is worth roughly ten dollars of enterprise value, to you, your partners, and everyone in the firm who holds a stake. It also means the reverse. Margin that depends on you personally, on undocumented process, on data you cannot produce on demand, gets discounted toward zero in any serious diligence.

The buyers only pay for what survives diligence without the founder in the room. **The institutionalization is not overhead on the way to the wealth. It is the wealth.**

8–12x

Durable fee-related earnings: the mid-market stake multiple, as of mid-2026¹

\$1 → \$10

A dollar of durable platform margin is roughly ten dollars of enterprise value

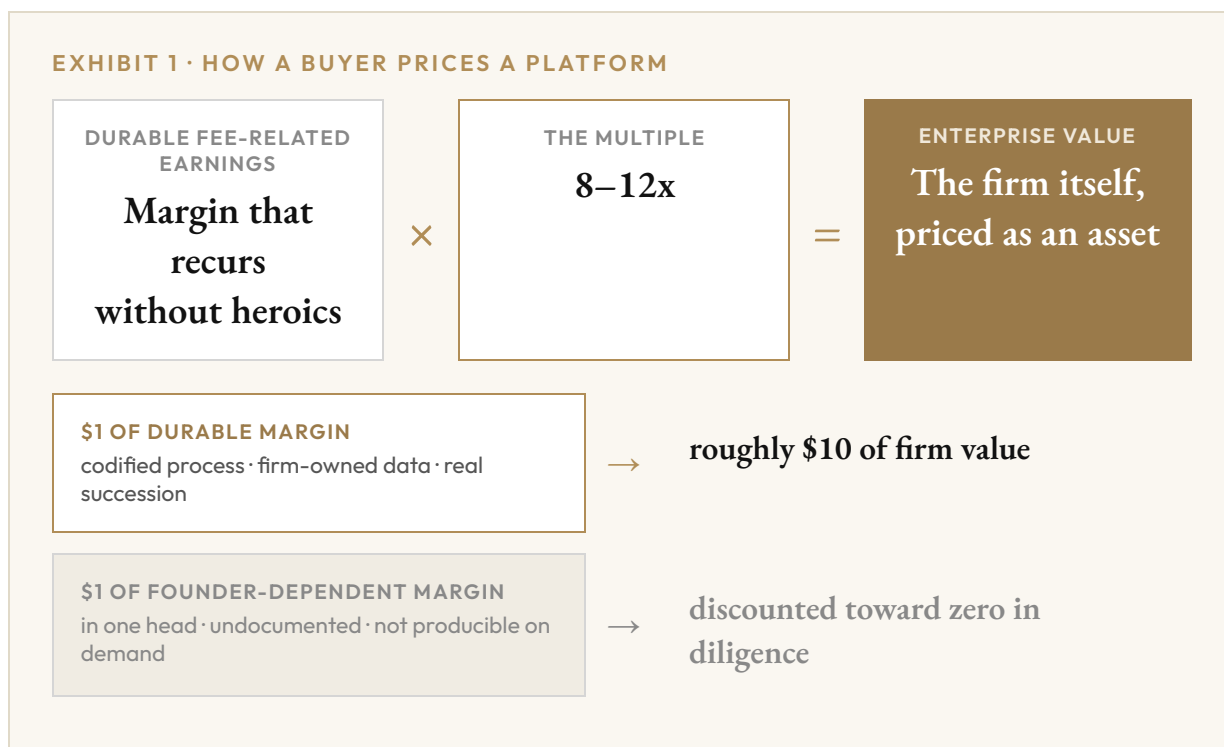
Eight figures

The typical gap between a firm today and the same firm at benchmark margin and durability

Most founders have never seen their number. The firms that have are not waiting for a capital event to start managing it. This playbook shows you how the number is built, where yours is leaking, and the seven-step order of operations that raises it, with AI entering exactly once, at Step 5, where it lands in fee-related earnings at the multiple instead of becoming another tax.

1. Mid-market real estate GP-stake pricing of roughly 8–12x durable fee-related earnings: Thesis Driven, "The GP Investment Landscape" (2026). Large-cap alternative managers trade meaningfully higher: 13–30x FRE, ~22x average (Goldman Sachs research, via Glasfunds, 2025). All figures as of mid-2026.

THE ARITHMETIC ON ONE PAGE



Everything else in this playbook is the working detail behind this exhibit: who the buyers are and why the band sits where it sits (Section 1), what happened to the firms that built for it and the firms that did not (Section 2), how the discount works (Section 3), how the earnings and the multiple are actually constructed (Section 4), the ten dimensions a diligence process scores (Section 5), and the seven steps that move the number (Section 6).

One page of arithmetic, honestly applied, changes how a founder reads every operational decision on his desk. The margin question stops being about this year's compensation pool. It becomes a balance-sheet question, at ten dollars to the dollar.

1

The Market That Prices You Now

A decade ago, the only price a mid-market real estate platform ever received was the one a strategic acquirer offered on the way to absorbing it. Today a specialist market prices platforms as going concerns, in minority-stake transactions that leave the founder in control. Understanding that market is the first step, because its arithmetic applies to your firm whether or not you ever sell a share.

SECTION 1 · THE BUYERS

A specialist class of capital, built to own pieces of firms like yours

GP-stakes investing began at the large-cap end, with Petershill, Dyal, and Blackstone buying into the biggest alternative managers. What matters for you is what happened next: the model moved down-market and specialized, and real estate got its own dedicated buyers.²

Buyer	Real estate footprint	What it signals
Almanac Realty Investors (Neuberger Berman)	\$8.9B invested across 59 real estate companies; \$800M partnership with NPS announced January 2025	The longest track record of platform (not fund) investing in real estate; their own vocabulary is “real estate platforms”
Bonaccord Capital Partners	Fund II closed at \$1.6B (January 2025); Fund III raising	Mid-market GP stakes as a scaled, repeatable strategy
Kudu Investment Management	Permanent-capital stakes across asset managers, including real estate specialists	The proof case in Section 2 is a Kudu deal
Cantilever Group	Targets managers with \$500M–\$5B AUM explicitly	A buyer whose stated target band is, precisely, firms your size
Hunter Point Capital	\$4.3B raised for GP financing solutions (June 2026)	Beyond stakes: credit against the platform’s own value
Others active	Investcorp (ISCP II, \$1.1B), RidgeLake (\$1.1B Fund I), Wafra, Blue Owl’s mid-market JV	Multiple bidders now exist at every size band

The market cleared roughly \$20 billion of GP-stakes transactions in 2025.³ The mid-market slice is smaller and growing faster, because that is where the un-institutionalized platforms, and therefore the value-creation opportunities, sit.

2. Buyer figures: almanacrealty.com and NPS partnership announcement (January 2025); Bonaccord Fund II close (GlobeNewswire, January 2025); cantilevergroup.com; Hunter Point Capital announcement (Businesswire, June 2026); Investcorp and RidgeLake fund announcements (2024–2025). All as of mid-2026.

3. Hogan Lovells market commentary citing With Intelligence data (~\$20B, 2025); PitchBook counted \$3.5B across 24 disclosed US PE GP-stakes deals in 2025, with real assets additional.

Why eight to twelve times, and why durable earnings only

The multiple is not sentiment. It rests on a finding that predates the GP-stakes boom: roughly two-thirds of the expected revenue of a private equity manager comes from fixed management fees, not carry.⁴ The fee stream is the franchise. Carry is real, but it is lumpy, years out, and contingent on exits, so a buyer discounts it heavily and sometimes to zero.

That is why the sell-side convention for valuing alternative managers, visible in public equity research for two decades, is a sum of the parts: fee-related earnings at a premium multiple, carried interest at a discounted one.⁵ Public large-cap managers earn 13 to 30 times FRE, roughly 22 times on average.⁶ Your mid-market band of 8 to 12 times is the same logic with two haircuts applied: less product diversification, and more founder dependence. The second haircut is the one you control, and it is the subject of Section 3.

What moves a firm inside the band

Within 8–12x, buyers price durability. Perpetual and open-end vehicles anchor the top of the band because their fees do not expire. Closed-end funds in their investment period price mid-band. Deal-by-deal and SMA-heavy revenue prices at the bottom, or below the band entirely. Re-up rates, margin stability, and whether the firm's data survives diligence without a founder narrating it move the number inside the band. Section 4 makes this mechanical.

Why this applies to you even if you never sell

Three reasons. First, the arithmetic prices your succession and your employee equity whether or not a buyer ever appears; a phantom-equity grant against a firm worth 3x thin earnings is a gesture, and against 10x durable earnings it is wealth. Second, LPs now underwrite the same durability the stake buyers do: 85 percent of institutional LPs report having rejected an investment over operational concerns alone, and 68 percent now weigh operational clarity above historical returns.⁷ Third, the option itself has value. A firm that could take a stake tomorrow negotiates everything else, lenders, JV partners, senior hires, from a different seat.

4. Metrick and Yasuda, "The Economics of Private Equity Funds," *Review of Financial Studies* (2010).

5. Sum-of-the-parts convention: sell-side research on Blackstone and peers (Morgan Stanley, Credit Suisse, Evercore ISI) has applied premium multiples to fee-related earnings and discounted carry since the early 2010s.

6. Goldman Sachs research via Glasfunds (2025).

7. CSC, survey of 150 institutional LPs (2025).

The window is open, and it is closing

Two signals from 2026 tell you where this market is heading, and both argue for moving before the crowd does.

Infrastructure became a headline use of proceeds. When Prime Finance took a Bonaccord stake in March 2026, the announcement led with technology and institutional infrastructure as intended uses of the capital.⁸ That is new. For the first decade of this market, stake proceeds were announced as fund commitments, product launches, and balance-sheet uses. The buyers have learned, from their own portfolios, that the operating layer is where the value creation actually happens. The next section shows the experiment that taught them.

The biggest platforms are institutionalizing innovation itself. Harrison Street, at \$109 billion, appointed the sector's first Global Chief Innovation Officer in June 2026.⁹ What a \$109 billion platform builds as a C-suite function today is what a \$2 billion platform will be asked about in diligence within a few years. The bar does not travel down-market slowly anymore.

Meanwhile, capital is concentrating

The fundraising backdrop sharpens the stakes. The top ten private equity funds captured 45.7 percent of all capital raised in 2025, up from 34.5 percent; experienced managers took 87.6 percent of commitments; and two-thirds of LPs plan to increase existing relationships rather than add new ones.¹⁰ Management fees across the industry have compressed to roughly 1.61 percent.¹¹ In that world, the mid-market platform's survivable position is durability: the firm that keeps its LPs, holds its margin, and can prove both. That position is buildable, and it is exactly what the stake market pays for.

Nothing here says sell. Most readers of this playbook should not, or not yet. It says: the price exists, the drivers are knowable, and every quarter you operate without managing to them is a quarter of compounding you do not get back.

8. Prime Finance and Bonaccord announcement (GlobeNewswire, March 2026).

9. Harrison Street announcement, harrisonst.com (June 2026); AUM per Commercial Observer interview (June 2026).

10. PitchBook (2025); Adams Street Global Investor Survey (2025).

11. Preqin (2025).

“

The buyers only pay for what survives diligence without the founder in the room. The institutionalization is not overhead on the way to the wealth. It is the wealth.

CHIRAG HATHIRAMANI · THE PLATFORM VALUE PLAYBOOK

2

The Natural Experiment

The public record now contains something close to a controlled experiment: platforms of similar size, in similar markets, that took the same kind of capital and made different choices about the operating layer. What happened next is the clearest evidence in this playbook, and none of it is mine. It is all public.

The firm that built the invisible middle

In December 2019, Pennybacker Capital, an Austin-based real assets manager then under two billion dollars, took a minority investment from Kudu. The announced uses were unglamorous: balance sheet, technology, staffing, growth.¹² No new office. No headline product launch. The interesting part is what the firm actually did with it.

It built what I call the invisible middle. A proprietary data platform, developed in-house, that turned operator feeds and asset performance into firm-owned truth. A genuine succession restructure at the C-suite, announced publicly in January 2023, not drafted in a drawer. A professional capital formation function, with a dedicated head hired in 2022. An evergreen vehicle. None of these show up in a deal announcement. All of them show up in diligence.

Then the visible results arrived, in the order the thesis predicts. The flagship fund went from \$510 million (2018) to over \$1 billion (2019) to \$1.6 billion (2024), one of the largest US real estate funds raised that year. Firm assets grew from under \$2 billion to roughly \$4.8 billion. In 2025 the firm raised over \$700 million for an entirely new asset class, and the announcement credited the data platform by name.¹³

Honesty requires two caveats, and they matter. No published attribution ties Pennybacker's returns to its platform, and its own fund announcements do not lead with it. The claim here is narrower and stronger: a mid-market firm took platform capital, spent it on the operating layer, and subsequently tripled its flagship while entering new products at institutional scale. The pattern is the argument. I had no role in any of it; I am citing a public record.

12. Kudu Investment Management and Pennybacker Capital announcement (December 16, 2019).

13. Fund and AUM figures: fund close announcements 2018–2024 (PRNewswire; PERE coverage of the 2024 close); infrastructure vehicle announcement (September 2025); succession announcement (January 2023). All public, as of mid-2026.

The firm that ran the visible playbook

Another manager in the same state took the same kind of capital three years later, from one of the most experienced platform investors in the market. On paper, the setups rhyme: comparable size, specialist thesis, respected founder, institutional minority partner with capital earmarked for growth.

What followed was the visible playbook: geographic expansion, new product lines, senior hires announced with appropriate fanfare. The parts of a platform you can photograph. What the public record does not show is the invisible middle. No data platform. No succession restructure. No industrialized capital formation function.

The public record does show the trajectory. The firm's perpetual vehicle grew by roughly \$26 million over three and a half years, effectively flat for a vehicle whose entire purpose is compounding. The next flagship fund, in market for over three years, has no announced close. Managing-director-level departures accumulated.¹⁴ Same state, same kind of capital event, opposite trajectory.

I am deliberately not naming this firm, because the point is not embarrassment. The point is that the difference between these two stories was not the sector thesis, the track record, or the quality of the capital partner. It was what the firm built underneath, in the years when building was optional.

The honest caution

The record contains a third case, and any advisor who omits it is selling you something. A New York manager took platform capital into a sector thesis that was breaking, and no amount of infrastructure could save the next fund.¹⁵ The system in this playbook multiplies a franchise. It does not resurrect one. If the market has left your thesis, fix that first; the operating layer comes second, and anyone who tells you otherwise is pricing their engagement, not your firm.

14. Trajectory from public filings and fund databases: SEC Form D amendments for the perpetual vehicle (2020–2026); fund-database status for the successor flagship; departure coverage in trade press. Firm anonymized deliberately.

15. Public record, 2018 transaction; successor fund filings show it never reached scale. Anonymized for the same reason.

SECTION 2 · THE LESSON

What the experiment actually isolates

Put the three cases side by side and the variable isolates itself. All three had capital. All three had track records good enough to attract a sophisticated institutional partner. Two had working sector theses. Only one built the invisible middle, and only that one compounded.

	The proof	The peer	The caution
Platform capital	2019, with technology and staffing as stated uses	Three years later, same capital class	2018
What got built	Data platform, succession, capital formation function, evergreen vehicle	New office, new product lines	Infrastructure could not matter
Sector thesis	Intact	Intact	Broken
Outcome	Flagship tripled; AUM ~2.5x; new asset class at scale	Perpetual vehicle flat; flagship unresolved; senior attrition	Next fund never scaled

Three implications for your firm. First, capital is not the constraint; two of the three firms show that capital without the build changes nothing. Second, the build precedes the results by years, which is why waiting for the capital event to start is already too late; the proof case began building the moment the ink dried, on capabilities that took three to four years to show up in fundraising. Third, the buyers have watched the same experiment. It is why infrastructure has started appearing as a headline use of proceeds, and why their diligence has learned to look for the invisible middle rather than the photographable one.

Every section that follows is a component of the invisible middle, priced.

Same state. Same kind of capital event. Opposite trajectories. The difference was not the sector thesis or the track record. It was the operating platform underneath.

SECTION 2 · THE NATURAL EXPERIMENT

3

The Founder Discount

Before a buyer opens your returns, they mark the whole firm down for one thing: how much of it walks out the door if you do. The discount is not an insult. It is a price on a risk you built into the firm, usually with the best intentions, one indispensable decision at a time. This section shows how it works, which means showing exactly how to remove it.

Your own fund documents price it first

You do not have to take my word that founder dependence is priced, because you have already signed the market's clearest statement of it. Most fund agreements carry a key-person clause: if the named founder steps away, the investment period suspends, capital calls stop, and new deals require investor consent. That is the market pricing founder dependence at its most literal, in your own LPA, negotiated by your own counsel.

A buyer of the platform reads the same clause and generalizes it. If the fund cannot operate without you, the fee stream the fund generates cannot be underwritten without you either. So diligence goes hunting for everything else that lives only in your head: the buy box, the LP relationship history, the reason the firm passed on the deals it passed on, the assumptions underneath the underwriting model. Whatever it finds there, it discounts.

The discount is not small, and it has two well-documented roots. The first is knowledge concentration: the standard finding in knowledge-management research is that roughly 42 percent of an organization's knowledge is unique to the individual who holds it, unrecoverable by colleagues if that person leaves.¹⁶ At professional-services turnover rates of around 13 percent a year,¹⁷ a firm that has not codified is not holding an asset; it is renting one, one resignation at a time. The second root is diligence behavior: operational failings alone now kill allocations. In one 2021 survey, 39 percent of institutional investors reported passing on funds that had already cleared investment diligence, over operations; by 2025, 85 percent reported having rejected investments on operational grounds.¹⁸

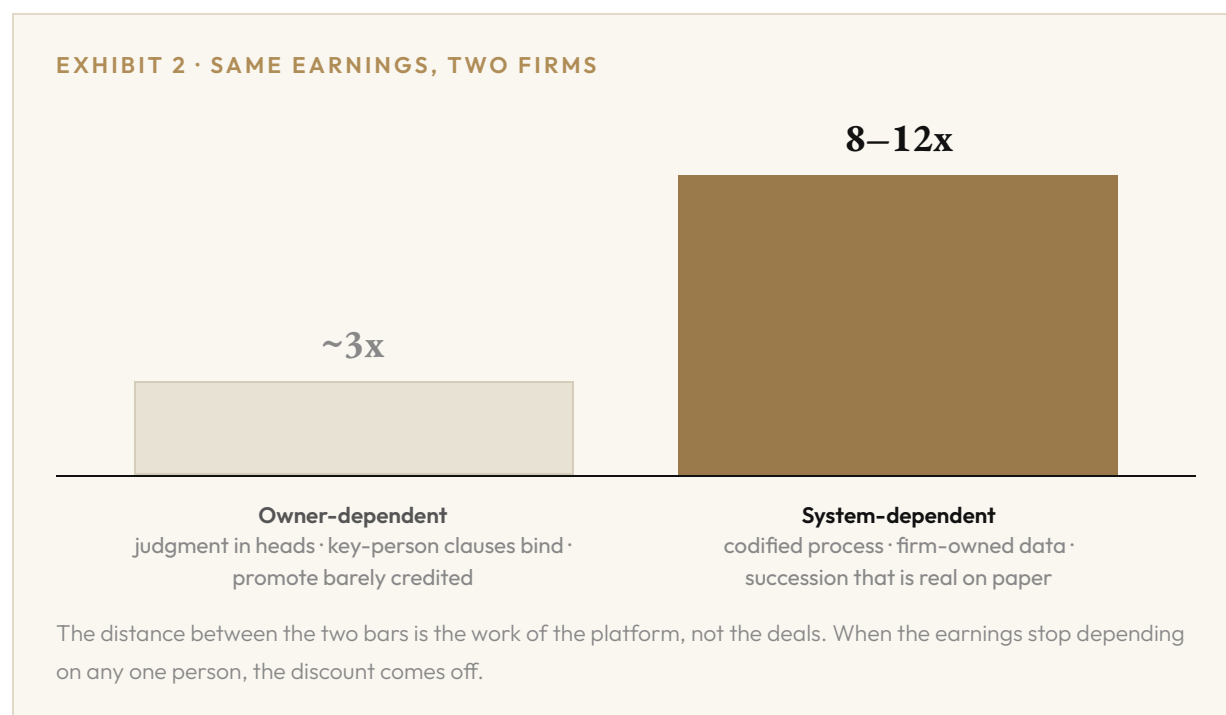
Founder dependence is the operational failing that contains the others. It is why the peer firm in Section 2 could add offices and products without adding a dollar of durable value: everything new still routed through the same two heads.

16. Panopto, Workplace Knowledge and Productivity Report (2018), still the standard reference.

17. Mercer, US Turnover Survey (2025).

18. CAIA investor survey (2021); CSC survey of 150 institutional LPs (2025).

SECTION 3 · THE SPREAD



Notice what the exhibit does not say. It does not say the system-dependent firm has better returns, better deals, or better people. Both bars are the same earnings. The spread is entirely a durability judgment, which means it is entirely operational, which means it is entirely buildable. A founder cannot decide to have a better track record. He can decide, this quarter, to start moving margin from the left bar to the right one.

And the carry deserves its own honest sentence. The promote is the number your partners toast at the holiday party, and a buyer pays almost nothing for it today: lumpy, years out, contingent on exits that may or may not come. This is not cynicism; it is discounting. The practical consequence is liberating. Since the promote barely prices, the entire lever is the fee stream, and the fee stream is the part of the business you control completely.

What removing the discount actually requires

The founder discount comes off when a diligence team spends two days inside your firm and cannot find the seam. Concretely, that means four tests, each of which is checkable from outside:

The firm can be described without your biography. The pitch, the process, and the edge survive translation into documents a stranger can read. If the honest answer to “what is your edge” begins with your name, the discount stays on.

Decision rights exist on paper and operate in practice. Who can approve what, to what dollar level, with what escalation path. Buyers read operating agreements before they read track records, because the operating agreement tells them whose exit breaks the firm.

The second layer owns something real. Succession that consists of a promise is priced as a promise. Succession that consists of named people with decision rights, economics, and a track record of running meetings you are not in is priced as durability. The proof case in Section 2 restructured its C-suite publicly, years before it needed to.

The firm’s knowledge is the firm’s. The buy box written and versioned. The pass decisions logged with reasons. The LP preferences in a system rather than a memory. This is Step 3 of the system, and it is where the 42 percent stops walking out the door.

None of this requires diminishing the founder. The firms that clear these tests are not less founder-led; they are founder-led on purpose instead of by default, with the founder’s judgment deployed on the calls that need it rather than embedded invisibly in every workflow. The deepest move a CEO makes is engineering his own dispensability into the firm’s design. That is what makes the firm worth anchoring, and worth a multiple.

4

Durability and the Multiple

This is the mechanical section: how the earnings number is constructed the way a buyer constructs it, what benchmark margin looks like, which durability drivers set your position inside the band, and a fully worked example you can rebuild for your own firm with a spreadsheet and an honest afternoon.

Constructing fee-related earnings the way a buyer does

Start with what recurs. Management fees on committed or invested capital, asset management fees, and the recurring service fees that arrive whether or not you close a deal this quarter. That is the revenue base. Now the buyer's adjustments, each of which surprises founders the first time:

Episodic fees get discounted or excluded. Acquisition, disposition, and financing fees swing with the deal calendar. A buyer treats them as cyclical, not durable, and haircuts accordingly. If half your revenue is transaction fees, half your revenue barely prices.

Costs get normalized. Above-market founder compensation gets added back; below-market gets deducted, because the buyer must assume a real cost of replacing you. One-time build costs come out. What remains is the platform's honest run-rate cost of operating.

The margin is what survives. Durable revenue minus normalized run-rate cost, expressed as a margin. In the mid-market, sub-30 percent margins read as sub-scale or over-founder-weighted; the benchmark for a well-run platform sits near 45 percent, and disciplined perpetual-vehicle managers run higher. The gap between your margin and benchmark is not shame; it is inventory. Every point of it is a dollar of potential FRE, at the multiple.

Then durability classifies it. The same dollar of FRE prices differently depending on what it is attached to. Margin attached to a perpetual vehicle with documented re-ups prices at the top of the band. Margin attached to a fund in year six with an unresolved successor prices at the bottom. Margin attached to a founder's personal deal flow prices as if it were already gone.

A useful discipline: build this construction for your own firm before anyone asks. The founders who find it uncomfortable are exactly the ones who benefit most, because every discomfort in the exercise is a discount a buyer would have applied silently.

SECTION 4 · THE DRIVERS

What sets your position inside the band

Within 8–12x, five drivers do most of the work. They are listed in rough order of weight, and every one of them is operational:

Driver	Top of band looks like	Bottom of band looks like
Fee durability	Perpetual and open-end vehicles; overlapping closed-end vintages	Deal-by-deal economics; SMA-heavy; single fund late in life
Founder independence	Firm passes the two-day diligence test without the founder narrating	Key persons everywhere; judgment undocumented
LP franchise	Documented re-up rates; institutional roster; reporting on cadence	Concentrated HNW base; every raise starts from zero
Data and process	Operator feeds, valuations, and asset data produced on demand from firm-owned systems	Data assembled by hand from operator PDFs when someone asks
Margin quality	Near benchmark, stable across quarters, produced by infrastructure	Thin, volatile, produced by heroics and deferred compensation

Two things to notice. First, track record is not on the list. It is the price of admission, not the differentiator; every firm a stake buyer looks at has returns worth looking at. Second, each driver is also an LP-diligence dimension. The same work that moves the multiple moves the next fundraise, which is why this playbook keeps insisting the arithmetic applies whether or not you ever sell: 53 percent of LPs now rank a manager's value-creation approach among their top selection criteria, right behind performance and team.¹⁹

The weakest driver governs. Buyers price the vulnerability, not the average, for the same reason lenders underwrite the covenant most likely to break. A firm with four drivers at benchmark and founder independence at zero is priced off the zero. Section 5 turns this into a scoring discipline.

¹⁹ McKinsey Global Private Markets Report (2026), survey of 300 institutional LPs.

SECTION 4 · WORKED EXAMPLE

A \$2 billion platform, priced end to end

A composite mid-market firm, with round numbers you can rebuild: \$2.0 billion of fee-earning AUM, blended fees just over 1.1 percent, margin in the mid-thirties, a vehicle mix of mostly closed-end funds with one modest perpetual sleeve, and operational maturity at the middle of every scale.

Step	Calculation	Result
Durable revenue	\$2.0B fee-earning AUM × ~1.125% blended	~\$22.5M
Fee-related earnings	\$22.5M × 35% margin	~\$7.9M
Durability position	Mid-band drivers; mixed vehicle durability	9–10x
Implied enterprise value	\$7.9M × 9–10x	~\$71–79M
Same firm at benchmark	\$22.5M × 45% margin × 11x	~\$111M
The gap	Benchmark value less today's midpoint	~\$37M

Read the last line slowly, because it is the whole playbook in one number. This firm is healthy. Its founder would describe it, accurately, as successful. And the distance between the firm as it runs today and the same firm at benchmark margin and durability is roughly thirty-seven million dollars of enterprise value, none of which requires a better market, a bigger fund, or a different track record.

That is the number I call the Platform Value Gap, and in the mid-market it is usually eight figures. Ten points of margin, at the multiple, is most of it; the durability re-rating inside the band is the rest. It is also, not coincidentally, the number that makes the build economics trivial: an operating-layer investment measured in hundreds of thousands, against a gap measured in tens of millions, is not a technology decision. It is a capital allocation decision, and capital allocation decisions get made.

SECTION 4 · DISCIPLINE

How to hold these numbers honestly

Valuation language earns trust only when it is conservative, so four disciplines govern every number in this playbook, and they should govern yours.

The multiple is not yours to promise. Nobody controls where a buyer prints. What you control are the drivers: durable FRE and documented durability. Manage the drivers; let the multiple be a consequence. Any advisor promising an exit multiple is describing their fee, not your firm.

The number is a floor concept, not a headline. The implied value in the worked example deliberately excludes carried interest, balance-sheet co-investments, and the option value of future funds. Real transactions price those too. Starting from the floor means every surprise in a real process runs in your favor.

Every figure carries an as-of date. The 8–12x band, the buyer roster, the LP-behavior statistics: all true as of mid-2026, all perishable. Re-verify quarterly anything you repeat to your board, and distrust any document about this market that is not dated.

The gap is credible only if the constraint is named. A gap number without a binding constraint is marketing. The honest version names the one thing that most limits the number today, margin, founder dependence, vehicle mix, data, and sequences the work against it. In my experience the binding constraint is rarely the one the founder would guess, which is precisely why pricing the platform is Step 1 of the system rather than a victory lap at the end.

~\$37M

The Platform Value Gap of a healthy, successful \$2 billion platform running at today's typical margin and durability. No better market required. No bigger fund. The same firm, at benchmark.

SECTION 4 · WORKED EXAMPLE, AS OF MID-2026

5

The Ten Dimensions Buyers Underwrite

Diligence is not a mood; it is a checklist. Whether the counterparty is a stake buyer, an institutional LP running operational due diligence, or a lender underwriting the management company, the same ten dimensions keep appearing. Score your firm against them honestly and you are reading your own diligence report before anyone commissions it.

SECTION 5 · DIMENSIONS 1-5 · THE ECONOMICS AND THE ENGINE

1 · Fee durability

What fraction of revenue arrives without a deal closing, and how long is it contracted to keep arriving? Vehicle structure dominates: perpetual and open-end fees compound; closed-end fees expire on a schedule everyone can read; episodic fees barely price.

2 · Margin quality

Not just the level but the character. Is margin produced by infrastructure, or by under-market compensation and heroic hours that a buyer must normalize away? A 40 percent margin that depends on the founder not paying himself is a 25 percent margin wearing a costume.

3 · Founder independence

The Section 3 tests: describable without the biography, decision rights on paper, a second layer that owns something real, knowledge in systems. The single heaviest dimension, because it multiplies every other one.

4 · Process codification

The high-stakes twenty percent of how the firm underwrites, decides, and manages, written by the people who do the work, versioned, and actually used. Binders that exist to pass audits score low; living documents that new hires learn from score high.

5 · Data infrastructure

Can the firm produce its own truth on demand? Operator feeds, valuations, asset performance, portfolio exposures, available in hours from firm-owned systems, rather than reassembled by hand from operator PDFs each quarter. This is the dimension diligence teams now test live, by asking and watching.

SECTION 5 · DIMENSIONS 6–10 · THE LEVERAGE AND THE FRANCHISE

6 · AI leverage

Whether the firm converts AI into fee-related earnings, or into scattered experiments and a verification burden. Note the dependency: this dimension cannot outscore dimensions 4 and 5, because AI compounds codified process and firm-owned data and multiplies chaos everywhere else. That gating is Step 5's entire argument.

7 · LP franchise

Re-up rates you can document, concentration you can defend, and a reporting cadence that would survive an unannounced visit. The franchise is measured in the second fund an LP commits to, not the first.

8 · Capital formation and deployment

Both pipelines as machines rather than personalities. On the raise side: pipeline discipline, diligence response times, a DDQ library that answers from documentation rather than memory. On the deploy side: sourcing-to-close throughput the firm can prove with conversion arithmetic. Durability shows here first, which is why buyers read both pipelines as a proxy for everything else.

9 · Product architecture

Whether the vehicle lineup is deliberate: flagship cadence, perpetual sleeve, successor products that extend the franchise rather than fragment the team. Product sprawl reads as opportunism; architecture reads as strategy.

10 · Governance

Decision rights, conflicts handling, compliance that operates rather than exists, and, increasingly, controls over systems that act: who authorized what, with what spend limits, and who can stop it within the hour. Diligence asks these questions about your people already; it now asks them about your software.

The weakest-link rule. Effective maturity equals your lowest dimension, not your average, because buyers price the vulnerability and LPs walk over single findings. A firm of nines with one three is, for pricing purposes, a three with excellent qualities. Inflate the flat tire first.

SECTION 5 · SCORING

Score it like the outsider will

Here is the pattern I see every time a leadership team scores itself: the founder's score arrives one to two full levels above where a structured outside scoring lands. Not from dishonesty. From the inside view: the founder scores the firm he knows he could run in an emergency; the outsider scores the firm that exists when nobody is performing.

Three corrections produce an honest score. **Score the artifact, not the intention.** If the buy box is “in your head, mostly,” the codification score is a 1, however good the buy box is. **Let the person who does the work score the dimension.** Your controller knows the real state of the data spine; your IR lead knows whether LP preferences live in a system or a memory. Their numbers will be lower than yours, and righter. **Ask the artifact test question for every dimension:** if a four-person diligence team arrived unannounced Wednesday morning, what could we hand them by Friday, and how much of it would need a founder in the room to explain?

Then convert the profile into the two outputs that matter. The lowest dimension is your binding constraint and your work plan. The overall profile, run through Section 4's arithmetic, is your implied value range and your gap. Together they answer the only two questions a board should ask about this whole subject: what is the number, and what single thing most limits it?

A self-serve version of this scoring exists online: the Platform Value Score, at chiraghathiramani.com/platform-value-score, runs the ten dimensions and the Section 4 arithmetic in about twelve minutes and returns an indicative value range and gap estimate. It is free, and this chapter is complete without it.

Effective maturity equals your lowest dimension, not your average. Buyers price the weakest link, and LPs walk over single findings.

SECTION 5 · THE WEAKEST-LINK RULE

6

The Seven Steps

The system, in the order the value equation demands. Phase One prices the platform. Phase Two builds the durability the multiple pays for. Phase Three multiplies. The order is not a preference; it is the physics of the arithmetic, and the most expensive mistakes in this market come from running it backwards.

EXHIBIT 3 · THE PLATFORM VALUE SYSTEM

PHASE ONE · PRICE

1 Price the Platform the implied value, the gap, the binding constraint

PHASE TWO · BUILD

2
**Remove the
Founder Discount**

3
Codify the Firm

4
**Build the Data
Spine**

5
**Compound the
Margin**
the gate: AI enters
only here

PHASE THREE · MULTIPLY

6
**Industrialize Capital Formation and
Deployment**
both pipelines: capital in, capital out, vehicles where
they meet

7
Earn the Multiple

STEP 1 · PRICE THE PLATFORM

Everything starts with the number

Step 1 is Sections 4 and 5 of this playbook, run on your firm: fee-related earnings constructed the way a buyer would construct them, margin and durability benchmarked against the mid-market, the implied enterprise value computed at market multiples, and the ten dimensions scored honestly. The output is a number most founders have never seen, the gap between the firm today and the same firm at benchmark, and one sentence naming the constraint that most limits it.

Why price first, when the building feels more urgent? Three reasons, each learned the expensive way.

The number sequences the work. Without it, operational improvement is a preference contest: the CFO wants the data spine, the founder wants the AI, someone read about EOS. With it, the sequence is arithmetic. If the binding constraint is founder dependence, no data project moves the number until Step 2 runs. Firms that skip the pricing build the wrong thing first with impressive energy.

The number recruits the coalition. An operating-layer program pitched as best practice competes with every deal on the pipeline for attention, and loses. The same program pitched as “thirty-seven million dollars of enterprise value, of which eleven million lands in the partners’ equity and the phantom pool” recruits the CFO, the COO, and every carry-holder in the firm as champions. Wealth is nearer and dearer than technology.

The number creates accountability. Re-scored annually, it turns institutionalization from a virtue into a managed metric: your implied value moved from \$75 million to \$89 million this year, because margin did this and durability did that. What gets priced gets built.

STEP 1 · IN PRACTICE

Running Step 1 on your own firm

The full exercise takes a competent CFO one honest week, most of it gathering rather than calculating:

- **Revenue decomposition.** Every fee stream tagged durable, semi-durable, or episodic, with its contractual life. The discomfort of this tagging is the diagnosis.
- **Cost normalization.** Founder compensation restated to replacement cost; one-time items removed; the honest run-rate margin computed.
- **Durability scoring.** The ten dimensions of Section 5, scored by the people who do the work, with the weakest-link rule applied without mercy.
- **The arithmetic.** FRE times the band position equals the implied range; benchmark margin and durability minus today equals the gap; the lowest dimension names the constraint.

Present the output on one page: the range, the gap, the constraint. Resist the appendix instinct. The board conversation you want is not about methodology; it is the capital allocation conversation the one-pager forces: should we spend a few hundred thousand dollars and two quarters to move the constraint one level, when the gap prices in the tens of millions?

A caution from experience: do not let the exercise become a valuation opinion. It is a management instrument, a compass rather than an appraisal. The moment it gets treated as a promise, conservatism dies and credibility with it. The language that keeps it honest: indicative, floor-concept, as-of-dated, drivers-not-outcomes.

Monday morning: ask your CFO for the revenue decomposition, tagged by durability, by Friday. That single artifact starts everything else, and its gaps will tell you more than this playbook can.

STEP 2 · REMOVE THE FOUNDER DISCOUNT

The firm stops being hostage to one calendar

Step 2 attacks the heaviest discount first. The work is governance, not technology, and it is the least glamorous and highest-priced work in the system: codified decision rights, real succession, and a second layer that owns something.

Decision rights, written and operating. Who approves what, to what dollar level, with what escalation. The test is not the document; it is whether the firm ran for a month, recently, with the founder approving only what the document says he approves. Most firms discover the document describes a firm they do not run.

Succession with economics. Named people, with decision rights and equity or phantom equity that vests against responsibility actually transferred. Succession without economics is an org chart; buyers read cap tables, not org charts. This is also where the coalition from Step 1 gets paid: the same institutionalization that removes the discount is what makes the second layer's equity worth holding.

The key-person audit. Read your own fund documents the way a buyer will. List every clause that names you, every consent that requires you, every relationship that would need a personal introduction to survive you. That list, prioritized, is Step 2's work plan. The goal is not to remove your name from the firm; it is to make each naming a choice rather than a load-bearing accident.

The proof case in Section 2 did this publicly and early, restructuring its C-suite years before any capital event required it. That ordering is the lesson: succession announced from strength reads as durability; succession announced under pressure reads as risk, and prices accordingly.

STEP 2 · IN PRACTICE

What changes, and what it is worth

The founder's week changes first. The calls that genuinely need founder judgment, pricing a marginal deal, reading an anchor LP, a senior hire, get more of it; everything else routes to owners with authority to decide. Founders report the same surprise: the firm does not wobble. It steadies, because decisions stop queueing behind one calendar.

The valuation change is the one this playbook can price. Recall the spread in Exhibit 2: owner-dependent earnings near three times; system-dependent earnings at eight to twelve. No other step moves a firm between those bars. Margin work (Steps 4 and 5) raises the earnings; only Step 2 changes what the earnings are worth per dollar.

Signals that Step 2 is real rather than performed:

- The firm holds its weekly leadership meeting with the founder absent, on cadence, and decisions ship.
- An LP's operational diligence interviews the team without the founder in the room, by the firm's own suggestion.
- The second layer kills a deal the founder liked, on documented criteria, and the criteria win.
- The founder takes ten consecutive working days away and the only artifacts are decision logs, not crises.

Each signal is checkable from outside, which is the point. Durability that cannot be observed cannot be priced.

Monday morning: run the key-person audit on your newest LPA. Count the clauses that name you. Every founder I have run this with guessed low.

STEP 3 · CODIFY THE FIRM

The high-stakes twenty percent becomes owned IP

Step 3 converts the firm's judgment from memory into property. Not everything, and this restraint is the step's discipline: the high-stakes twenty percent of how the firm actually underwrites, decides, and manages, written by the people who do the work, versioned like the intellectual property it is.

What belongs in the twenty percent: the buy box, with the reasons behind each boundary; the underwriting assumptions library, including the market-by-market rules of thumb that live in senior heads; the IC memo standard and the decision log, including passes and the reasons for them; the LP preference record; the asset management playbooks for the situations that recur. What does not belong: the exhaustive process manual nobody reads. A comprehensive binder is not an asset; it is a warning sign that documentation was performed rather than used.

Why this prices. Recall the two numbers from Section 3: roughly 42 percent of organizational knowledge is unique to the person holding it, and professionals turn over at roughly 13 percent a year. Compound those and an uncoded firm sheds irreplaceable knowledge continuously, invisibly, forever. A buyer cannot price what might leave in a resignation. A firm whose buy box is a versioned document with five years of decision history attached has converted that leak into an asset that survives anyone's departure, including yours.

Codification is also the precondition for everything downstream. The data spine (Step 4) needs to know what the firm considers true; the AI layer (Step 5) can only be taught judgment that has been written down. Firms that skip to the technology discover they have automated ambiguity.

STEP 3 · IN PRACTICE

Written by the doers, kept alive by use

The failure mode of codification is the consultant-written binder that dies on a shelf, so the working rules run opposite to instinct:

- **The doer writes; the senior corrects.** The analyst drafts the underwriting standard from what actually happens; the MD marks it up. The markup is where the judgment transfers, and, not incidentally, where the next generation learns.
- **One artifact at a time, provoked by use.** Codify the buy box the week a live deal tests it; the IC standard the week a memo disappoints. Documents written for real decisions get used for the next one.
- **Version and date everything.** A buy box that shows its own evolution, what the firm believed in 2022 versus what it believes now, and why the change, is more valuable than a perfect current-state document, because it proves the firm learns on paper.
- **Retire aggressively.** Anything unconsulted in two quarters gets merged or killed. The corpus stays the twenty percent.

The observable payoff arrives in diligence, in a form buyers increasingly test live: ask the newest associate to walk through how the firm screens a deal, without a senior in the room. In the codified firm, the associate answers from the same corpus the MD uses, and the diligence team writes the sentence every founder wants in the report: the process is the firm's, not any individual's.

Monday morning: ask your most senior deal person to write the buy box, one page, boundaries and reasons. Then ask your newest analyst to mark what surprised them. The distance between those two documents is your codification backlog, ranked.

STEP 4 · BUILD THE DATA SPINE

Firm-owned truth, produced on demand

Step 4 is the least visible step and the one diligence tests hardest: the firm's numbers, operator feeds, valuations, asset performance, portfolio exposures, flowing into firm-owned systems that can produce truth on demand, rather than reassembled by hand each quarter from whatever format the property managers sent.

The spine has four layers, and the order is the architecture. A master data model first: standardized definitions of NOI, occupancy, and the other twenty numbers the firm lives on, because a firm with three definitions of NOI has zero. Connectors second: automated feeds from property management and accounting systems, replacing the monthly copy-paste liturgy. Validation third: the rules that catch economic-versus-physical occupancy confusions and reclassified expense lines before they reach an IC memo or an LP letter. Only then the analytics and reporting layer everyone wanted to build first.

Why the spine prices so heavily: it is the dimension a diligence team can verify in an afternoon, live. Ask for a variance analysis on a named property. Ask for portfolio exposure by market, by vintage, by covenant. Watch what happens. The firm that produces answers from systems in minutes has proven, in the most credible way possible, that its reported numbers are the firm's numbers. The firm that says "we'll get back to you" has just repriced itself, whatever its returns.

The spine is also where LP diligence and buyer diligence converge: both run the same playbook, because operational due diligence taught them to. Building it once answers both audiences for years.

STEP 4 · IN PRACTICE

Ninety days to the first honest dashboard

The spine is not a two-year IT program, and treating it as one is how it dies. The working sequence, repeatable at mid-market scale:

- **Weeks 1–3: the definitions fight.** One meeting, the right people, and a signed one-page data dictionary. This is governance wearing engineering's clothes; have the fight once, on purpose.
- **Weeks 3–8: two feeds, not twelve.** Connect the two systems that carry the most weight, usually property management and fund accounting, and let the remaining feeds earn their sequence by pain.
- **Weeks 6–10: validation rules from the last three mistakes.** Every firm has a story: the reclassified expense line, the occupancy definition that slipped. Encode those first; they are validation rules with proof of relevance.
- **Weeks 10–13: one dashboard someone already needed.** Not the executive cockpit; the variance report the asset manager assembles by hand every month. Ship the thing that deletes a manual liturgy, and the team's trust arrives with it.

Two honesty notes. First, data work only stays funded when it is priced; this is why the spine sits inside a valuation program rather than an IT budget, where it would compete with laptops. Second, the spine is the gate for Step 5, and the gate is absolute. AI on top of unvalidated data does not create margin. It creates fluent, confident, wrong answers about your own portfolio, at scale.

Monday morning: ask for last quarter's portfolio occupancy, by property, with the definition attached. Time how long it takes and count the hands it passes through. That is your spine baseline, measured.

STEP 5 · COMPOUND THE MARGIN

Only now does AI enter, and this gate is the brand

Step 5 is where artificial intelligence joins the system, and its position is the system's most important design decision. AI enters after codification and after the data spine, because AI is a compounding mechanism, and compounding requires something true to compound. Layered on codified process and firm-owned data, AI converts capacity into margin. Layered on chaos, it manufactures a verification burden that costs more than the work it replaced.

What Step 5 actually deploys, at mid-market scale and in production rather than pilot: first-pass underwriting assembled from the firm's own standards and comp history, so analysts stress-test rather than assemble; deal screening against the codified buy box, so the pipeline arrives ranked; quarterly reporting drafted from the spine's validated data in the firm's approved language, so a two-week liturgy becomes a two-day review; variance detection that reads every property every month and surfaces the three that need a human.

Notice what every item has in common: each lands directly in fee-related earnings, by deleting hours from the cost line or adding capacity without headcount, and each depends explicitly on Steps 3 and 4 existing first. The screening needs a written buy box. The reporting needs validated data. The pattern is general: **AI's output quality is a function of the codified judgment and owned data underneath it, which is why the invisible middle has to be built before it can be compounded.**

This is also where the margin math from Section 4 becomes concrete. A mid-market platform's path from a mid-thirties margin to benchmark is mostly recovered coordination: the assembly, reconciliation, and re-formatting hours that consume expensive people. Step 5 is how those hours come back, and at the multiple, every recovered durable dollar is roughly ten dollars of firm value.

STEP 5 · IN PRACTICE

Production discipline, not pilot theater

The industry's AI record is honest and bad: the large surveys keep finding that most firms' initiatives stall in pilots, and only a small fraction reach goals.²⁰ The failures share a shape, and the discipline that avoids them fits on one page:

- **One workflow, proven, then the next.** A single scout with a single broken process beats a firm-wide mandate every time, because visible proof recruits the next adopter and mandates recruit resistance.
- **Verification is architecture, not vigilance.** Every automated output carries its sources, and a checking layer verifies against the spine before a human sees it. The human reviews judgment, not arithmetic.
- **Permissions are earned rung by rung.** Systems draft before they send; recommend before they decide; and money movement and LP communication keep a named human signature permanently. These are the controls an operational diligence team will ask about by name: attribution, spend limits, and who can stop the system within the hour.
- **Measure cost per shipped outcome.** Not seats, not experiments: what a completed underwriting, a shipped report, a screened pipeline costs before and after. Margin is the only scoreboard that survives a board meeting.

Readers who want the AI layer at depth should take the companion volume: the CRE AI Playbook, free at chiraghathiramani.com, is this step expanded to fifty pages, the technology deep dive inside the value framework you are holding. The order of the two documents is the argument: platform value first, AI inside it.

20. Representative findings as of mid-2026: JLL research on real estate AI pilots (2025: 88% of firms began pilots; ~5% achieved goals); MIT State of AI in Business (2025: ~95% of enterprise generative AI pilots without measurable P&L impact); Deloitte CRE Outlook (2026).

Two pipelines, one growth engine

Phase Three multiplies what Phase Two built, and it starts from what a real estate company actually is: a firm that raises capital, places it well, and executes what it underwrote. Step 6 industrializes the first two as a single machine, because in a fund manager they are one flywheel. Fee-earning AUM, the driver of everything the multiple prices, is the product of capital raised and capital deployed; a firm that raises what it cannot place grows nothing a buyer pays for.

The raise side. The machine exists apart from the founder's charisma: a pipeline with stages and coverage discipline; re-up management as a year-round function, because the re-up rate is the single most legible durability statistic a firm produces; a DDQ and diligence-response library that answers from documentation in days; reporting cadence kept as a trust signal, quarter after quarter. The context makes this urgent: two-thirds of LPs now prefer increasing existing commitments over adding managers, and experienced managers captured 87.6 percent of 2025 commitments. The cheapest capital you will ever raise is the re-up.

The deploy side. The sourcing-to-close machine, running on everything Phase Two built: the codified buy box screening the market, the firm's comp and market data underneath the underwriting, and the Step 5 throughput turning offering memoranda into IC-ready packages in hours. What the buyer reads: the firm can place the next fund as efficiently as it placed the last, and prove it with pipeline arithmetic rather than anecdotes.

Where the pipelines meet: product. Vehicles convert two-sided capacity into sticky fee-related earnings: the perpetual sleeve that anchors the top of the band, the successor-fund cadence that compounds the franchise, products that extend the platform rather than fragment the team. The proof case in Section 2 launched an evergreen vehicle and then an entirely new asset class on exactly this machinery.

One boundary, held deliberately: this step industrializes the throughput and discipline of your investing, never the judgment inside it. Deal selection stays yours. A machine that promises alpha is a machine selling something.

STEP 6 · IN PRACTICE

Measured like pipelines, run like a product

The observable state of an industrialized Step 6, checkable in any quarter, on both sides of the machine:

- **The re-up rate is a managed number.** Documented by fund, decomposed by reason on every non-re-up, reviewed by leadership on cadence. If the firm cannot state it without a research project, Step 6 has not started.
- **Diligence response time is measured in days.** The standard DDQ answered from the library inside a week; the data room assembled from the spine inside two.
- **Reporting ships on the same day each quarter.** Unglamorous, and among the most-cited operational facts in allocator references. Cadence is character, observed.
- **The deal pipeline has coverage arithmetic.** Opportunities screened per period, screen-to-LOI and LOI-to-close conversion, time-to-close, all trended. The numbers prove deployment capacity the way re-ups prove revenue durability.
- **Deployment pace matches the raise.** Dry-powder aging tracked and honest: capital raised against a working machine, not a hope. A widening gap between the two pipelines is the earliest warning either one is broken.
- **Every closed deal feeds the machine.** Outcomes flow back into the buy box and the assumptions library, so the dealcraft compounds the way the fundraise does.

The second-order effect holds on both sides: founder hours concentrate where judgment moves outcomes, the anchor conversations and the deals worth fighting for, while the machine carries cadence and assembly. The firm raises more and places more with less founder time, which is itself the durability the buyer is underwriting.

Monday morning: ask for two numbers, in writing: your re-up rate across the last two funds, and your screen-to-close conversion over the trailing year. Whichever takes longer than a day is the finding.

STEP 7 · EARN THE MULTIPLE

Optionality is the point

Step 7 assembles what the first six built into a firm that could transact, on the founder's terms, whether or not it ever does. The components are unglamorous and decisive:

Clean platform financials. The management company's own P&L, produced to the same standard as fund reporting, with the durable and episodic streams separated and the normalization arguments already made. Most mid-market platforms cannot produce this today; every buyer starts by asking for it. A firm that hands it over in a day has set the tone for the entire process.

Event readiness, rehearsed. The Section 7 dry run, executed before any counterparty exists: the data room generated rather than assembled, the diligence questions answered from documentation, the findings fixed on your own calendar. The vehicle architecture that converts capacity into durable fees was built back in Step 6; here it gets presented for what it is, strategy legible on one page.

The documented equity story. Not a sales document; the firm's own account of what it is, what compounds, and why the next decade's economics are more durable than the last's, supported by the spine's data rather than adjectives.

Hold all three and every door is open: a minority stake at a multiple that did not exist for the firm two years earlier; a succession funded by the firm's own value rather than the founder's savings; employee equity that means something because it is priced against a real number; or deliberate independence, chosen from strength. The system's end state is not a sale. It is a firm whose founder decides, rather than a firm whose circumstances do.

STEP 7 · AND THE ORDER, ONCE MORE

Why the sequence is the system

Run the seven steps backwards and each one fails predictably. AI before the spine automates ambiguity. The spine before codification digitizes three competing definitions of NOI. Capital formation before founder-independence industrializes a bottleneck. A transaction before any of it sells at the discount, permanently. The order is not pedagogy; it is dependency.

Run forward, each step funds the next. The pricing recruits the coalition. Governance unlocks codification, because people document what they own. Codification feeds the spine; the spine gates the AI; the AI pays for everything by landing in margin; and the margin, made durable and legible, multiplies through both pipelines into the multiple.

EXHIBIT 4 · THE THREE JOBS, MAPPED

What a real estate company must do well	Where it lives in the system
Raise capital efficiently	Step 6, the raise side: re-ups, diligence response, cadence, channels
Find and close great deals	Steps 3–5 build the dealcraft (buy box, data, throughput); Step 6 industrializes deployment
Execute the business plans	Steps 3–5: codified process, the data spine, the margin engine

And the value itself: Step 1 prices it, Step 2 makes it durable, Step 7 monetizes it.

Timeline honesty: the full arc runs in quarters and years, not weeks; the proof case took roughly three years from capital to visible fundraising results. But the first artifacts, the priced gap, the key-person audit, the buy box, the first honest dashboard, arrive inside a quarter, and each is individually valuable. And the one-sentence version, for the elevator: **price the platform first, build the durability the multiple pays for, then multiply.**

7

Stake-Readiness: The Buyer's Checklist, Run Early

If a capital event sits anywhere on your horizon, two to five years counts, there is a discipline that changes the outcome more than any banker selection: run the buyer's diligence on yourself, before any buyer arrives, while there is still time to fix what they will find.

SECTION 7 · THE CHECKLIST

What a stake buyer's diligence actually requests

The lists converge across buyers, because the risks converge. The recurring requests, organized by what they are really testing:

Request	What it tests
Management company P&L, three years, with normalization detail	Whether durable earnings exist and the firm knows its own number
Fund documents: key-person, consent, and transfer provisions	How much of the firm's value is hostage to one calendar
Operating agreements and equity/phantom schedules	Whose exit breaks the firm; whether the second layer owns anything
LP schedule with re-up history and concentration	Revenue durability, evidenced rather than asserted
Live data requests: variance analysis, exposure cuts, decision trails	Whether the spine exists; whether numbers are the firm's or a founder's narration
Process corpus: buy box, IC standards, decision logs	Whether the judgment is owned IP or walking inventory
Compliance, conflicts, governance of systems that act	Whether the firm operates its policies or merely has them

Every row maps to a section of this playbook, which is the deeper point: stake-readiness is not a separate project. It is the Platform Value System, audited. A firm that has run Steps 1 through 6 walks into diligence with the data room generated rather than assembled, and the process becomes a confirmation exercise instead of an excavation.

Findings priced on your calendar, not theirs

The economics of early diligence are asymmetric in your favor. A finding surfaced two years before a process is a work plan: the succession gap gets closed, the vehicle mix gets rebalanced, the margin normalizes, and none of it appears in anyone's price. The same finding surfaced by a buyer's diligence team is a negotiation: priced at the moment of maximum leverage against you, or worse, it becomes the reason a term sheet quietly does not arrive.

Buy-side practice prices findings by severity; sell-side discipline is making sure the severe ones no longer exist. Hence the sequence for a founder with a two-to-five-year horizon:

- **Now:** the self-diagnostic (Section 8), the Step 1 pricing, and the key-person audit. One quarter, mostly internal.
- **Year one:** the binding constraint, worked. For most firms this is Step 2 or Step 4; for margin-thin firms, Step 5's recovered capacity.
- **Year two:** a full dry run against the checklist opposite, ideally with an outside reader who has seen real diligence, because the inside view cannot see its own seams.
- **When the event is real:** process-fit becomes its own decision, which buyer archetype, what governance terms, what the capital is for, and it deserves specialist advice beyond this playbook's scope.

One more honest sentence about incentives: the sell-side advisors who run these processes enter at the transaction and are paid on it closing. Nobody in that value chain is paid to build your asset in the years before, which is precisely why the building, done early and on your own calendar, captures value that a process, however well run, can only distribute.

A finding surfaced two years early is a work plan. The same finding surfaced by a buyer is a price adjustment.

SECTION 7 · STAKE-READINESS

8

The Self-Diagnostic

Twelve questions, answerable in one honest hour with your CFO. Score each 1 (no, and it would embarrass us in diligence) to 5 (yes, documented, and operating). No question asks about returns; every question is something a buyer or an LP's operational diligence will test.

#	Question	Dimension
1	If no new deal closed this year, how much of our revenue would still arrive, and could we show that from documents rather than from memory?	Fee durability
2	If we paid the founder what it would actually cost to replace him, would the firm still have a healthy margin, or has his paycheck been quietly subsidizing it?	Margin quality
3	If a diligence team spent two days here and the founder stayed out of the room, could the team get their answers, or does everything route back to one person?	Founder independence
4	Have we named the people who would carry the firm after the founder, and do they hold real authority and real ownership today, not a promise of it?	Founder independence
5	Is the way we decide what to buy written down somewhere a new analyst could learn from, including the deals we passed on and why?	Codification
6	If someone asked for portfolio occupancy by property this afternoon, could we produce it within the hour, with everyone agreeing on what occupancy means?	Data spine
7	When AI or automation produces work for us, does it arrive showing where its numbers came from, with a named person accountable before it leaves the firm?	AI leverage
8	Do we know what share of our investors came back for the next fund, and could we explain, one by one, why the others did not?	LP franchise
9	Could we answer a standard due diligence questionnaire in a week from documents we already have, and could we show how many deals we screen for each one we close?	Capital formation & deployment
10	Could we explain on one page why our funds and vehicles are structured the way they are, or has the lineup accumulated one opportunity at a time?	Product architecture
11	If someone asked us to prove our compliance and conflicts policies are actually followed, could we show them evidence, or only the policy document?	Governance
12	Have we ever honestly priced the firm itself, and could we name the one thing that most limits that number today?	The whole system

SECTION 8 · READING YOUR SCORE

The lowest number is the finding

Do not average. Averages flatter, and diligence does not. Your effective score is your lowest answer, and the question that produced it is your binding constraint, the same one Step 1 would have named with a spreadsheet.

Mostly 4s and 5s, lowest a 3: you are closer to the top of the band than most of your peers, and the remaining work is legibility: documenting what already operates so an outsider can price it. Section 7's dry run is your highest-value next move.

A profile of 3s: the typical honest mid-market result, and roughly the worked example on page 23: a healthy firm carrying an eight-figure gap. The sequence is the system's: price it properly, then work the constraint. Expect the constraint to surprise you.

Any 1s on questions 3 through 6: founder dependence or data. These price heaviest and take longest, which makes them first regardless of what feels urgent. The good news buried in the low score: these are also the steps with the steepest value slope; a firm moving founder independence from 1 to 3 does more for its multiple than any margin program.

Question 12 at a 1 is the most common answer in this industry and the cheapest to fix. An honest week of CFO time produces the number, and the number changes the conversation about everything else.

Whatever the profile, resist the instinct to launch six initiatives. The system moves one constraint at a time, on purpose: focus is what makes the compounding visible, and visible compounding is what keeps a leadership team funding the unglamorous work.

The online version, the Platform Value Score at chiraghathiramani.com/platform-value-score, runs a fuller thirty-statement scoring with the valuation arithmetic attached, free, in about twelve minutes. This chapter is complete without it; the online version adds the computed value range and gap.

A CLOSING NOTE

The firm is the product

Twenty years on the principal side taught me a small number of durable things, and this playbook is built on the least glamorous of them: the deals are what the firm makes, and the firm is what the founder makes. Most of us spent our careers optimizing the first and improvising the second. The market has now put a price on the difference.

I do not believe every reader should raise platform capital, and I do not believe infrastructure saves a broken thesis; Section 2's caution case is in this document on purpose. I do believe, from watching the natural experiment run in public, that the mid-market firms which build the invisible middle in the next few years will compound away from the ones that do not, in fundraising, in talent, and in the value of the firm itself, and that the window in which this is a differentiator rather than a requirement is measured in years, not decades.

Everything in these pages can be run by your own team. The arithmetic is public, the checklists are complete, and the Monday-morning moves are deliberately small. If you run them and want company for the harder middle stretch, the back page says how to find me. Either way, run them. The number exists whether or not you look at it, and the founders who look first keep the difference.

The first step is a number. Find out what your firm is worth.

SOURCES AND METHODOLOGY

References

All market figures are current as of mid-2026 and will age; the mechanisms they illustrate are the point.

Figures worth repeating to a board should be re-verified quarterly. Primary sources by section:

Pricing and market structure (Sections 1, 4). Mid-market real estate GP-stake pricing of ~8–12x durable FRE: Thesis Driven, “The GP Investment Landscape.” Large-cap comparison (13–30x, ~22x average): Goldman Sachs research via Glasfunds (2025). Market volume: With Intelligence via Hogan Lovells (~\$20B, 2025); PitchBook (24 disclosed US PE deals, \$3.5B, 2025). Buyer facts: company announcements and websites, Almanac/NPS (January 2025), Bonaccord Fund II (January 2025), Cantilever, Hunter Point (June 2026), Investcorp, RidgeLake. Fee economics: Metrick & Yasuda, *Review of Financial Studies* (2010). Sum-of-the-parts convention: published sell-side research on listed alternative managers (Morgan Stanley, Credit Suisse, Evercore ISI).

The natural experiment (Section 2). Pennybacker Capital: Kudu transaction announcement (December 16, 2019); fund closes 2018–2024 (PRNewswire; PERE); succession announcement (January 2023); capital formation appointment (2022); infrastructure vehicle (September 2025). Peer and caution cases: SEC Form D filing trails, fund-database status, and trade press; anonymized deliberately, and no non-public information is used anywhere in this document.

Diligence behavior and knowledge risk (Sections 3, 5). CSC survey of 150 institutional LPs (2025: 85% / 68%); CAIA survey (2021: 39%); McKinsey Global Private Markets Report (2026: 53% of 300 LPs); Panopto Workplace Knowledge Report (2018: 42%); Mercer US Turnover Survey (2025: ~13%).

Industry context (Sections 1, 6). Capital concentration: PitchBook (2025: 45.7% / 87.6%); Adams Street Global Investor Survey (2025). Fee compression: Preqin (2025: 1.61%). AI adoption base rates: JLL (2025), MIT State of AI in Business (2025), Deloitte CRE Outlook (2026). Timing signals: Prime Finance/Bonaccord announcement (March 2026); Harrison Street announcement (June 2026).

Methodology. The worked example (page 23) is a composite at stated round numbers, not a client; rebuild it with your own inputs. The valuation frame is a floor concept: it excludes carried interest, balance-sheet assets, and future-fund option value, and no statement in this playbook is a valuation opinion or an offer of one. The frameworks are documented at book length in *The Platform CEO* (theplatformceo.com).

Chirag Hathiramani

Creator of The Platform Value System · Former Chief Investment Officer

Author of *The Platform CEO* · Austin, Texas

The first step is a number. The Platform Value Score prices your firm across the ten dimensions buyers underwrite and returns your indicative enterprise value range, free, in about twelve minutes.

chiraghathiramani.com/platform-value-score

chi@chiraghathiramani.com

chiraghathiramani.com · theplatformceo.com

linkedin.com/in/addrealvalue