The CRE AI Playbook

How middle-market real estate PE firms build the operational infrastructure to scale like institutions — and compound an advantage no competitor can replicate.

Chirag Hathiramani | Real Estate Consulting

Second Edition · Q2 2026

Contents

01

Executive Summary

The thesis, the data, and what this playbook delivers.

The firms at the top of the PERE 100 did not get there by picking better deals than everyone else. They got there by building operating systems that let them deploy capital at scale without proportional headcount growth. The platform is the product. The deals flow through it.

This distinction has always mattered. What has changed is the cost of building institutional infrastructure. Components that used to require dedicated departments of fifteen people can now be built as automated workflows at a fraction of the cost. A well-designed system lets a twelve-person team run a platform that used to require fifty. The question is no longer whether middle-market CRE PE firms can build institutional-grade infrastructure. It is whether they will.

80%

of AI projects fail to reach meaningful production deployment — twice the failure rate of non-AI IT projects.[2]

88%

of CRE investors have started piloting AI — but only 5% report achieving most of their program goals.[4]

The macro environment has made this urgent. Bain & Company's 2026 Global Private Equity Report documents a structural shift in deal math: a decade ago, a typical buyout required roughly 5% annual EBITDA growth to hit target returns. Today, that number is closer to 12%.[18] The tailwinds that powered the industry's growth — low rates, rising multiples, cheap debt — are gone. What remains is a hypercompetitive market where operational excellence is the primary lever for generating returns.

Who this is for

This playbook is written for CEOs, CIOs, and senior operators at CRE PE firms with $500 million to $5 billion in AUM. It is not a technology manual. It is a framework for thinking about operational infrastructure — with AI as the enabling layer, not the objective.

How to use this document

A CEO with four minutes can scan the stat callouts and section titles and absorb the thesis. A CIO with an hour can read it cover to cover and walk away with a specific diagnostic framework and implementation roadmap to discuss with their team on Monday. An LP conducting operational due diligence can evaluate whether the GP who shared this document thinks institutionally about their platform.

Every claim in this playbook is footnoted to its primary source. Every framework is grounded in what the technology can actually do as of mid-2026, not what it might do someday. Every example is either drawn from public sources with attribution or explicitly labeled as a hypothetical scenario.

To understand why this infrastructure matters now, we need to look clearly at where the industry actually stands.

> *The winning firms will build systems, not slogans. They will invest in talent and AI, and move from full potential diligence to execution on Day 1.*

— Hugh MacArthur, Chairman, Global Private Equity Practice, Bain & Company, 2026

02

The Current State of CRE PE Operations

Where the industry actually stands — the manual processes, the fragmented systems, and the structural gap between how firms operate and what institutional capital now demands.

The manual process reality

Most CRE PE firms between $500 million and $5 billion in AUM operate with a technology architecture that is, at its core, a collection of spreadsheets held together by email. Deal pipeline tracking lives in a shared Excel workbook or a lightly used CRM. Underwriting models are built from scratch for each transaction or inherited from the last analyst who held the seat. Monthly property financials arrive from property managers in different formats across the portfolio. Quarterly LP reports are assembled over three weeks by pulling data from multiple systems, reconciling discrepancies manually, and formatting the output in PowerPoint or Word.

None of this is a failure of intelligence or effort. It is the natural consequence of how these firms grew. Most were founded by operators who built their reputations on deal judgment and relationships — not on systems architecture. The lean team and the hands-on principal who touches every deal are signature strengths of an entrepreneurial firm. But somewhere between $500 million and $1 billion in AUM, they become ceilings.

The CRE technology stack problem

The CRE industry's core technology vendors — Yardi, MRI Software, RealPage, Argus, AppFolio, Juniper Square, Dealpath — were each designed to solve a specific problem well. None were designed to talk to each other. The result is an archipelago of data: property financials in one system, deal pipeline in another, investor records in a third, underwriting models in Excel, and institutional knowledge in email archives and people's heads.

JLL's 2025 Global Real Estate Technology Survey of more than 1,500 senior decision-makers found that while 88% of CRE investors have started piloting AI, more than 60% remain strategically, organizationally, and technically unprepared to execute on their ambitions.[25] The money is flowing — 87% are increasing their technology budgets specifically for AI — but outcomes are lagging behind investment. Only 5% of firms report achieving most of their AI program goals.[4]

> **60%+**
>
> **of CRE investors remain strategically, organizationally, and technically unprepared to execute on their AI ambitions — despite 87% increasing tech budgets specifically for AI.[25]**

The fundraising pressure

The pressure from the capital markets side is equally acute. Buyout fundraising dropped 16% in 2025, and the number of funds closing fell for the fourth consecutive year.[31] Distributions to LPs as a percentage of NAV have been mired below 15% for four straight years — a record for the modern private equity industry.[10] More than half of LPs report being limited in making new commitments because prior commitments have not been drawn down.[32]

In this environment, LPs are increasingly selective. The ILPA LP Sentiment Survey 2025–26 found that more than half of LPs believe they have more leverage with GPs than they did twelve months earlier.[20] The median GP now offers 33 cents of no-fee coinvestment per dollar of fee-bearing capital — a 25% effective reduction in revenue.[11] The economics are

compressing from both sides: the cost of generating alpha is rising while revenue per dollar of AUM is declining.

Among the most effective responses to rising costs and compressing fees is operational leverage. And increasingly, the most practical path to building that leverage at the speed the market demands is infrastructure that scales without proportional headcount growth.

The shift from financial engineering to operational value creation

Bain's 2026 report frames this as a fundamental inflection point for private equity. The industry is experiencing a K-shaped recovery: a subset of elite funds is on the upswing while everyone else is working harder for diminishing results. The sorting mechanism is infrastructure. The funds raising capital fastest have both strong returns and strong distributions. The funds struggling have weak DPI, no systematic edge, and no data-backed story for LPs.[27]

The McKinsey Global Private Markets Report 2026 arrives at the same conclusion: outcomes will increasingly be shaped by deliberate operational choices — not market dynamics alone.[12] Alpha is less likely to emerge from leverage and multiple expansion. It will come from firms that can systematize sourcing, accelerate underwriting, monitor assets continuously, and demonstrate operational rigor to an LP base that has become far more demanding.

This operational gap has a specific cost — one that most firms have accepted as normal. The next section puts a number on it.

03

The Coordination Tax

The hidden cost that every CRE PE firm pays — and most have accepted as normal.

There is a cost embedded in every CRE PE firm's operations that never shows up on the income statement. I call it the Coordination Tax: the aggregate time, effort, and opportunity cost consumed by finding data, reconciling formats, chasing approvals, and reassembling information that already exists somewhere in the firm.

The Coordination Tax is not a technology problem. It is a systems architecture problem. It is the cost of operating without a platform.

Where senior time actually goes

At most mid-market firms, sixty to seventy percent of team time goes to data assembly, formatting, monitoring, and coordination. Thirty to forty percent goes to the judgment, relationships, and decision-making that actually creates value. The firm is paying senior talent to do junior work — not because anyone wants to, but because no system exists to separate the two.

Consider how a typical quarterly LP report gets produced. At a lean mid-market firm, the asset manager pulls property-level financials and prepares the numbers and narrative in conjunction with the accountant. The IR person then assembles and formats the quarterly report. At larger shops with dedicated reporting staff, the handoffs multiply: one person pulls the data, another reconciles against the GL, a third formats, a fourth writes commentary. Regardless of the firm's size, the CIO reviews the output, catches discrepancies that reflect different source data rather than analytical errors, sends it back for correction, and reviews again. The entire cycle consumes 400 to 500 person-hours annually at a typical mid-market

firm. Eighty percent of that effort is data gathering and formatting. Twenty percent is the analysis and narrative that LPs actually read.

The math

The Coordination Tax is quantifiable. Consider a ten-person firm with $4 million in annual fully-loaded G&A — C-suite compensation alone often exceeds $1.5 million before you account for VP, Associate, and Analyst salaries, benefits, office, technology, data subscriptions, travel, and professional services. The blended fully-loaded hourly cost is roughly $200. If the team works 20,000 hours per year and 60% is coordination — finding data, reformatting it, reconciling discrepancies, chasing approvals — that is 12,000 hours at $200 per hour. That is $2.4 million per year spent on operational friction that creates no value for LPs, generates no returns, and builds no competitive advantage.

> **$2.4M**
>
> Annual Coordination Tax at a typical 10-person CRE PE firm — the cost of operational friction that creates zero value for LPs. Even a 25% reduction recovers $600,000 in productive capacity annually.

Even a 25% reduction through structural remediation recovers $600,000 in capacity annually — the equivalent of two to three junior hires worth of productive time redirected from formatting spreadsheets to the work that actually drives returns.

The Coordination Tax across the operating model

Function	Where Coordination Tax Hides	Typical Hours/Year
Quarterly LP Reporting	Data collection from multiple systems, reconciliation, formatting, multi-round review cycles	400–500
Deal Screening	Manual parsing of broker emails, entering data into trackers, comp assembly from disparate sources	300–400
Underwriting	Model construction from scratch, assumption sourcing, manual sensitivity analysis	200–300 per deal
DDQ Responses	Searching prior responses, assembling attachments, reconciling data across submissions	150–250
Asset Management	Collecting PM reports, re-entering data, building budget-to-actual manually	500–800
IC Memo Assembly	Pulling data from underwriting, market research, legal, and structuring into memo format	40–80 per deal

How to identify it in your own firm

Track one week in detail. For every task, ask two questions: Is this task finding, formatting, or reconciling data? And could a system do this if the data were structured and connected? The gap between what your team spends time on and what your team should spend time on is your Coordination Tax. Your honest answer is your growth ceiling.

> *The infrastructure that got you here is exactly what prevents you from getting there. That is not a criticism. It is a diagnosis — and it represents one of your greatest leverage points for generating a competitive advantage.*

The connection to the $1M-per-head firm

There is a metric that clarifies the stakes: one million dollars of revenue per head. If your firm generates $8 million in fee revenue with twenty people, you are at $400,000 per head. The same $8 million with eight people puts you at $1 million per head. Same revenue. Completely different firm economics, margin structure, and LP narrative.

The Coordination Tax is the primary reason most mid-market firms are stuck at $400,000 per head. The team is not underperforming. The infrastructure is forcing senior talent to spend the majority of their time on work that does not require their judgment. Reducing the Coordination Tax does not require replacing people. It requires building the systems that let them do the work they were hired to do.

Reducing it starts with knowing where your firm stands. The next section provides the diagnostic framework — ten dimensions that determine whether AI creates value or creates more work.

> *Most of us built our firms on hustle and relationships. That is not a weakness. It is just not an operating system.*

04

The AI Maturity Spectrum

Where your firm sits today, where the leaders are, and the ten dimensions that determine whether AI creates value or creates more work.

Three levels of AI integration

Production-grade AI implementation is not a single capability. It is three distinct levels of integration, each building on the last. Most firms stop at Level 1 and wonder why the results are not compounding.

Level 1: Individual Tool Use

An analyst uses AI to draft a memo, pull comps, or format a model. The gains are real — individual productivity jumps measurably. But they stay isolated. Nobody has connected these individual wins into the investment process. This level represents about twenty percent of the available value.

Level 2: Firm Intelligence

The AI system operates with your firm's specific context: underwriting criteria, target market assumptions, lease abstraction templates, LP reporting formats, IC memo structure, historical deal data. Without this context, AI produces generic output that sounds authoritative but is not calibrated to your standards. With it, the system produces work product that reflects your firm's philosophy and risk appetite. This is where most of the uncaptured value lives.

Level 3: Strategic Alignment

The AI's output lands in the right place, at the right time, in the right format, with the right validation checks. The system is woven into workflows. Deal screening feeds underwriting

feeds IC feeds asset management feeds reporting. Every step informs the next. This is what turns scattered AI wins into a compounding platform.

McKinsey's 2025 State of AI survey confirms this pattern at the enterprise level: 88% of organizations use AI in at least one function, but only 6% qualify as "high performers" where AI contributes meaningfully to EBIT.[24] The gap between Level 1 and Level 3 is not a technology gap. It is an infrastructure gap.

The AI Maturity Index: ten dimensions

The AI Maturity Index (AIM Index) evaluates a CRE PE firm across ten weighted dimensions that span the full operating model — deploying capital, raising capital, and running the platform. Each dimension is scored across five maturity levels, from Tribal (undocumented, hero-dependent) to Agent-Native (AI as a team member with defined responsibilities and human oversight at decisions).

INTERACTIVE TOOL

Take the AI Maturity Index

Score your firm across all ten dimensions in three minutes. The AIM Index generates a composite maturity profile and identifies your highest-leverage improvement areas.

chiraghathiramani.com/aim-index

Dimension	Weight	What It Measures
Deal Sourcing & Origination	15%	Pipeline systematization, channel diversity, screening infrastructure
Underwriting & IC Process	15%	Model standardization, assumption libraries, IC memo consistency
Asset Management & Value Creation	12%	Property-level execution, BP tracking, operational reporting
Portfolio Intelligence & Risk	12%	Portfolio-wide analytics, concentration risk, scenario modeling
Capital Markets	10%	Fundraising pipeline, debt origination, LP prospecting
Investor Relations & Reporting	10%	LP reporting cadence, DDQ response, transparency
Institutional Knowledge	8%	Knowledge capture, decision logs, data architecture, system integration
Team & Talent Readiness	8%	AI fluency, training, workflow ownership, hiring evolution
Strategic Governance	6%	Board-level AI strategy, budget allocation, competitive positioning
Compliance & Data Security	4%	AI governance policy, audit trails, cybersecurity baseline

Where most firms sit today

In diagnostic work across the $500 million to $5 billion range, the median firm scores at Level 2 (Documented) across most dimensions and Level 1 (Tribal) on two or three. The binding

constraint is almost always Institutional Knowledge — the pillar nobody thinks about and the one that determines whether everything else works.

The binding constraint rule applies: a firm's effective maturity is its lowest-scoring dimension, not its average. A firm at Level 4 on seven dimensions but Level 1 on Institutional Knowledge has an effective readiness of Level 1. AI cannot call the person who knows how things work.

Level 3

The minimum viable maturity for production AI deployment. Most mid-market CRE PE firms are at Level 2 or below. The goal is not Level 5 everywhere — it is Level 3 across all ten dimensions.

Five maturity levels

Level	Description	What AI Can Do
1: Tribal	Processes live in people's heads. Hero-dependent. No documentation.	Nothing meaningful. AI will hallucinate or create more work than it saves.
2: Documented	Some templates and SOPs exist. Not enforced. Inconsistent adoption.	Basic task automation. High error rate on firm-specific context.
3: Standardized	Enforced standards. Templates used. Workflows repeatable.	Reliable automation: deal screening, report generation, variance analysis.
4: Integrated	Systems connected. Data flows automatically. Performance measured.	Multi-step workflows: end-to-end deal processing, automated LP reporting.
5: Agent-Native	AI as team member with defined responsibilities, access controls, human oversight at decisions.	Autonomous execution with human-in-the-loop at judgment calls.

Reaching Level 3 across all ten dimensions requires a specific kind of infrastructure underneath — a foundation that most firms have never explicitly designed. That foundation is the Data Spine.

05

The CRE Data Spine

The four-layer infrastructure model that determines whether your AI investments produce institutional-quality output or expensive noise.

Every AI failure in CRE PE traces back to the same root cause: the firm tried to build the application layer without building the infrastructure underneath it. They bought the tool before building the platform.

The CRE Data Spine is a four-layer model that describes what institutional-grade data infrastructure actually looks like in a real estate portfolio. Most firms try to build Layer 4 without Layers 1 through 3. That is why their dashboards lie and their AI output requires manual verification that costs more time than the original process.

> **43%**
>
> **of organizations cite data quality and readiness as their top obstacle to AI success — ahead of model sophistication, talent, or budget.**[14]

Layer 1: Master property data model

Standardized definitions for NOI, occupancy, rent roll format, CapEx categories, and chart of accounts across the entire portfolio. This sounds mundane. It is the single most important investment a firm can make in its data infrastructure.

The problem is specific to CRE: when one property manager reports "occupancy" as economic occupancy and another reports physical occupancy, and both numbers flow into the same portfolio dashboard, every downstream analysis is contaminated. When one property uses cash-basis NOI and another uses accrual, the portfolio-level aggregation is meaningless. When CapEx is categorized differently across assets, comparing value-add execution across the portfolio is impossible.

Layer 1 is a one-page data dictionary that defines every metric the firm tracks. It costs nothing but time and discipline. Most firms skip it because it is not interesting work.

Layer 2: Automated connectors

Automated data feeds from property management systems (Yardi, MRI, RealPage, AppFolio), accounting platforms, and third-party data sources into a centralized reporting layer. The key word is automated. If someone is downloading a CSV, reformatting it, and uploading it into another system, that is not Layer 2. That is the Coordination Tax in disguise.

JLL's chief technology officer put it precisely: you cannot have an AI strategy without a data strategy.[34] Layer 2 is where the data strategy becomes operational.

Layer 3: Validation rules

Automated checks that catch when data does not conform to the master data model. Economic versus physical occupancy mismatch. GAAP versus cash NOI inconsistency. Rent roll totals that do not reconcile with GL data. CapEx draws that exceed budget without an approved change order.

Layer 3 is where most firms discover how much of their reported data has been unreliable. The discovery is uncomfortable and essential. Without it, every AI application built on top produces outputs that look authoritative and are quietly wrong.

Layer 4: Monitoring, analytics, and AI

This is the layer everyone wants to build first. Portfolio dashboards. Automated variance analysis. AI-generated LP reports. Deal screening agents. Underwriting copilots.

Layer 4 only works when Layers 1 through 3 are in place. The AI is only as good as the data it consumes. An underwriting copilot trained on inconsistent rent roll formats will produce inconsistent outputs. A portfolio monitoring dashboard built on un-validated property financials will surface false signals. A deal screening agent without standardized intake criteria will rank deals against undefined standards.

McKinsey's research confirms this at the enterprise level: organizations that achieve significant financial returns from AI are twice as likely to have redesigned their data workflows before selecting modeling techniques.[16]

> *You cannot have an AI strategy without a data strategy. The companies that prioritized building out their data platforms are the ones picking up steam.*
>
> — Yao Morin, Chief Technology Officer, JLL

What "clean data" actually means in CRE

Clean data in CRE PE does not mean perfect data. It means data that is consistent enough across the portfolio that automated systems can consume it without human translation. Specifically, it means consistent property identifiers across all systems, a standardized chart of accounts that every property manager follows, rent rolls in the same format with the same field definitions, CapEx categories that match budget line items, and financial reporting on the same basis (cash or accrual, consistently applied).

The firms that get this right do not have better technology. They have better data discipline. And that discipline compounds: every quarter of clean data makes the next quarter's analytics more reliable, the AI outputs more trustworthy, and the LP reports more credible.

With the Data Spine in place, AI can operate effectively across the full investment lifecycle. The next section maps where the impact is greatest — and what maturity level each stage requires.

> *No LP ever invested in a fund because the GP had a Yardi license. What LPs invest in is the platform — the operational discipline, the systematic process, the institutional rigor that the tool enables.*

06

Impact Across the CRE PE Value Chain

Where AI creates measurable leverage at each stage of the investment lifecycle — the specific pain points, the mechanisms, and the honest assessment of what requires human judgment.

The CRE PE value chain runs from deal sourcing through disposition. At each stage, there are tasks that AI can execute with high reliability, tasks that benefit from AI-assisted analysis with human oversight, and tasks where human judgment remains essential. The firms getting results know which is which. The firms wasting money do not.

Deal Sourcing & Market Mapping

The pain point

Deal flow at most mid-market firms depends on a principal's network. Brokers email offerings. The acquisitions team manually reviews each one. Promising deals get logged in a tracker — sometimes. The pipeline is invisible to anyone who does not sit in the daily deal flow. If the head of acquisitions leaves, the sourcing channel narrows overnight.

How AI addresses it

Automated screening of incoming deal flow against defined criteria — target markets, asset class, deal size, return thresholds. Systematic monitoring of listing platforms for off-market indicators. Comp assembly from multiple data sources consolidated into a single view. The mechanism is pattern matching at scale, not deal judgment. AI screens. Humans decide.

Maturity level required:

Level 2 (Documented criteria). **ROI impact: $$$**

Underwriting & Screening

The pain point

Every analyst builds models with slightly different assumptions, from slightly different sources, in slightly different formats. When a model breaks, only the person who built it can debug it. IC packages take days to assemble because the underwriting output has to be manually translated into the memo format.

How AI addresses it

Standardized model templates with assumption libraries that carry source citations and vintage dates. Automated sensitivity analysis across defined scenarios. Draft IC memos generated from underwriting outputs in the firm's established format. The time from deal identification to IC-ready package compresses from days to hours — not because the analysis is faster, but because the assembly is automated.

Maturity level required:

Level 3 (Standardized). **ROI impact: $$$$**

Due Diligence & Execution

The pain point

Due diligence is document-intensive and time-pressured. Lease abstracts, environmental reports, title exceptions, survey reviews, estoppels — each requires extraction, analysis, and synthesis under a compressed timeline. The coordination overhead between legal, environmental, engineering, and the investment team is enormous.

How AI addresses it

Document extraction and abstraction at scale. Lease abstraction that identifies non-standard clauses and flags deviations from the LOI. PSA versus closing statement discrepancy

identification. The mechanism is pattern recognition on structured documents — precisely the kind of task where AI reliability is highest and the cost of manual execution is most acute.

Maturity level required:

Level 3 (Standardized). **ROI impact: $$$**

Asset Management & Hold Period Operations

The pain point

Business plan execution monitoring is quarterly at most firms. The asset manager receives monthly financials from the property manager, re-enters data into the firm's tracking system, manually calculates budget-to-actual variances, and produces a narrative update. By the time a variance is identified, investigated, and escalated, weeks have passed. Construction draw tracking often lives in a separate system entirely.

How AI addresses it

Continuous monitoring replaces periodic reporting. Automated budget-to-actual variance analysis with threshold-based alerts. Lease expiration tracking with tenant credit monitoring. CapEx variance identification against approved budgets. The shift is from retrospective reporting to real-time awareness — the asset manager learns about a problem when it happens, not when the quarterly report reveals it.

Maturity level required:

Level 3–4 (Standardized to Integrated). **ROI impact: $$$$**

LP Reporting & Investor Relations

The pain point

Quarterly LP reports consume 400 to 500 person-hours annually at a typical mid-market firm. DDQ responses require a scavenger hunt across prior submissions, firm documents, and team members' memories. Metrics are calculated differently depending on who builds the report. The IR team spends more time on production than on relationships.

How AI addresses it

Automated report assembly from standardized data pipelines. DDQ response generation from a maintained answer library. Consistent metric methodology applied across every reporting cycle. The IR team's role shifts from production to relationship management — from assembling reports to interpreting them for LPs and using the freed capacity for fundraising strategy and co-investment structuring.

Maturity level required:

Level 3 (Standardized). **ROI impact: $$$**

Capital Formation

The pain point

Most mid-market firms manage fundraising the same way they manage deal flow — through relationships and manual tracking. The LP pipeline lives in the principal's head or a lightly maintained spreadsheet. There is no systematic follow-up cadence, no CRM-driven touchpoint tracking, and no data-backed segmentation of the investor base.

How AI addresses it

CRM intelligence that tracks LP engagement patterns, identifies re-up timing, and flags relationship cooling. Automated meeting preparation briefs that surface an LP's prior questions, allocation history, and co-investment preferences. Distribution notice production from GL data through waterfall calculation to compliance check. The capital formation process becomes as systematic as the deal process.

Maturity level required:

Level 3–4 (Standardized to Integrated). **ROI impact: $$$**

Disposition & Exit

The pain point

Exit timing decisions rely heavily on market intuition. The data that would inform a rigorous hold-versus-sell analysis — comparable transaction velocity, cap rate trends, remaining business plan upside, LP liquidity preferences, fund-level return attribution — exists across multiple systems and requires manual assembly to evaluate.

How AI addresses it

Scenario modeling that synthesizes property-level performance, market conditions, and fund-level return implications. Automated comparable transaction monitoring. Draft disposition memos assembled from asset management data. The judgment on when to sell remains human. The analytical foundation for that judgment becomes systematic.

Maturity level required:

Level 4 (Integrated). **ROI impact: $$**

> *The AI is invisible — the mechanism, not the message. Every sentence a GP delivers to an LP about operational sophistication can be enabled by AI. None of them should mention it.*

Deploying AI across these seven stages raises an immediate governance question: which decisions should AI touch, which should it support, and which must remain entirely human?

07

The Tiered Trust Framework

Which decisions AI should touch, which it should support, and which must remain entirely human — mapped to the CRE PE operating model.

Most CRE PE firms skip the trust definition entirely. They either give AI no authority — so nobody uses it — or no guardrails — so the IC does not trust the output. The governance framework is the missing piece, and it is not a technology problem. It is a decision-rights problem.

McKinsey's 2026 AI Trust Maturity Survey found that nearly two-thirds of organizations cite security and risk concerns as the top barrier to scaling agentic AI — well ahead of regulatory uncertainty or technical limitations.[8] The average responsible AI maturity score across industries is just 2.3 out of 4, with only one-third of organizations reporting governance maturity at Level 3 or higher.[21]

Three tiers of AI autonomy

Tier	Autonomy Level	CRE PE Applications
Tier 1: Full Autonomy	AI executes end-to-end. Human reviews exceptions only.	CRM data entry, listing alert screening, rent comp assembly, file naming, data formatting, meeting transcript summarization
Tier 2: Human-in-the-Loop	AI produces output. Human reviews before action.	Underwriting sensitivity analysis, tenant credit assessment, LP report drafts, DDQ response assembly, IC memo first drafts, variance analysis narratives
Tier 3: AI-Assisted Only	AI provides analysis. Human makes and owns the decision.	IC investment recommendations, disposition timing, capital call decisions, LP communication strategy, fund structure decisions, hiring

The principle underneath the framework is simple: AI drafts, humans decide. AI accelerates the path to a decision. It does not make the decision. Every AI-generated output requires human review before any communication to an LP, lender, or legal counterparty; any financial figure in an IC memo or investor communication; any recommendation affecting investment strategy or capital allocation; and any regulatory filing.

The minimum viable AI governance policy

If an LP asked you today for your AI data governance policy, what would you hand them? If the answer is nothing, this framework provides the minimum viable version.

Zone	Data Types	AI Policy
GREEN	Public market data, published research, general CRE knowledge, public listings	May be processed by any AI system without restriction.
YELLOW	Internal deal memos, anonymized models, aggregated portfolio data, draft LP communications	Enterprise-grade AI only (zero-data-retention). No investor names, addresses, or counterparty identities.
RED	Investor PII (names, SSNs, bank details, K-1s), unredacted legal documents, capital account statements	Never processed by any external AI system. Internal systems only. Violations require immediate reporting.

Why governance is a competitive advantage

ILPA is developing supplementary DDQ modules specifically for real estate, private credit, and infrastructure — adding technology governance questions to the standard diligence process.[30] Firms that can hand an LP a documented AI governance policy, a clear data classification framework, and an audit trail for AI-assisted outputs will have a structural advantage in fundraising over those who improvise answers during due diligence.

With governance defined, the next question is practical: what should your firm build internally, what should it purchase, and what expertise should it bring in through advisory?

> *There is a specific kind of frustration that comes with knowing your team is capable of so much more, but the infrastructure will not let them show it.*

08

Build, Buy, or Borrow

The CRE PE firm's decision framework for AI infrastructure — and the third option that most mid-market firms actually need.

The standard framing — build versus buy — misses a third option that matters enormously for CRE PE firms in the $500 million to $5 billion range.

Most firms in this segment cannot justify $1 million or more annually in fixed engineering headcount, and they cannot wait 12 to 18 months for a custom build. But they also cannot simply purchase a SaaS tool and expect it to work without workflow redesign. The orchestration layer — the system that connects individual tools into an integrated platform — requires domain expertise that neither the vendor nor the internal team typically has.

What to build (proprietary)

Firm-specific scoring models trained on your deal history. Underwriting assumption libraries calibrated to your markets and investment philosophy. IC memo templates that reflect your firm's analytical framework. Knowledge bases built from your investment committee's decision history. These are proprietary by definition. No vendor can provide them. They are the components of the Compounding Loop — the advantage that deepens with every transaction.

What to buy (commodity)

CRM (Juniper Square, HubSpot). Property management platforms (Yardi, MRI, AppFolio). Deal management (Dealpath). Document extraction tools. Accounting platforms. These are commodity infrastructure — necessary, but not differentiating. The selection criteria should be

integration capability, not feature count. The system that connects well to other systems is more valuable than the system with the most features in isolation.

What to borrow (advisory)

The orchestration layer. The governance framework. The implementation roadmap. The data spine architecture. The diagnostic that identifies which pillar to fix first. These require deep understanding of both the technology landscape and the CRE PE operating model. They are best delivered as advisory engagements that transfer capability to the firm over time — not as permanent dependencies.

This is the fractional Chief AI Officer model: strategic oversight of AI deployment, vendor evaluation, team training, and operational architecture at a fraction of a full-time hire. The engagement builds the firm's internal capability rather than creating an ongoing consulting dependency.

The CRE vendor landscape

Category	Key Platforms	Integration Reality
Property Management	Yardi, MRI Software, RealPage, AppFolio, Greystar, Asset Living, RPM Living	APIs available but inconsistent. Data extraction often requires manual configuration per PM.
Deal Pipeline & CRM	Dealpath, Intapp DealCloud, Dynamo Software, Affinity, Navatar (Salesforce-native)	Strongest integration ecosystems in mid-market CRE PE. Dealpath and DealCloud lead for deal-specific workflows.
Investor Portals & IR	Juniper Square, Allvue, AppFolio Investment Management, FIS Digital Data Exchange, Altus Group	Growing automation for LP reporting and capital account management. Juniper Square dominates the mid-market.
Fund Accounting	FundCount, Allvue, FIS Private Capital Suite (Investran), LemonEdge	Critical for waterfall calculations and K-1 generation. Integration with IR portals varies.
Fund Administration	SS&C Technologies, Citco, Apex Group, Alter Domus, IQ-EQ, NAV Fund Administration	Outsourced NAV, investor services, and regulatory reporting. Scale-dependent cost structure.
Underwriting	Argus, proprietary Excel models	Argus has limited API access. Most firms rely on Excel with no version control.
Market Data & Analytics	CoStar / Real Capital Analytics, PitchBook, Green Street, MSCI, Bloomberg, Cherre	Strong APIs. The connective tissue between property-level and market-level data.

Category	Key Platforms	Integration Reality
Compliance & Regulatory	ACA Group, Kroll, IQ-EQ, Waystone	AIFMD, FATCA, Form PF, and Schedule K-1 compliance. Critical for institutional LP requirements.
Document Intelligence	Horizontal AI tools (Claude, GPT), Orbital	Lease abstraction and due diligence document processing rapidly maturing.

Even with the right build-buy-borrow decisions and a clear vendor landscape, most AI initiatives still fail. Understanding why — and specifically how these failures manifest in CRE PE — is essential to avoiding the same fate.

09

Why Most CRE AI Initiatives Fail

The four failure modes — each observed in production deployments — and the structural reasons mid-market firms get stuck between pilot and platform.

The RAND Corporation's landmark study — based on structured interviews with 65 experienced data scientists and engineers — identified five root causes of AI project failure. More than 80% of AI projects fail to reach meaningful production deployment, exactly twice the rate of non-AI IT projects.[2] The root causes are organizational, not technical: misunderstood problem definition, inadequate data, technology-first mentality, insufficient infrastructure, and problem difficulty.

In CRE PE, these universal failure modes take specific shapes.

Failure Mode 1: Data quality and integration

This is the most common failure mode and the most CRE-specific. The rent rolls come in different formats. The NOI definitions are inconsistent. The property management data does not reconcile with the GL. The AI tool is given fragmented, inconsistent inputs and produces fragmented, inconsistent outputs. The team blames the AI. The problem is Layer 1 of the Data Spine.

Gartner predicted that 60% of AI projects unsupported by AI-ready data would be abandoned through 2026.[23] In CRE PE, the percentage may be higher because the data fragmentation problem is structural — it is baked into the vendor ecosystem.

Failure Mode 2: Governance and fiduciary constraints

This is the most CRE-specific failure mode. Private equity firms operate under fiduciary obligations to their LPs. An AI-generated IC memo that contains an inaccurate statement is not merely a software error — it is a potential fiduciary breach. An AI system that processes LP capital account data through an external API without proper data governance creates regulatory exposure. The stakes are categorically different from deploying AI in a software company's customer support function.

The firms that stall at this point do so because they have no framework for which decisions AI can touch and which it cannot. The Tiered Trust Framework in Section 7 exists to solve this specific problem.

Failure Mode 3: Talent and organizational readiness

This is the most underestimated failure mode. AI implementation requires someone in the organization who understands both the technology and the business well enough to scope projects, define success criteria, evaluate outputs, and manage the ongoing system. Most mid-market CRE PE firms have neither a technologist who understands real estate nor a real estate professional who understands technology architecture.

McKinsey's 2025 survey found that only 1% of executives describe their gen AI rollouts as "mature."[33] The maturity gap is not about the tools. It is about the organizational capacity to absorb them.

Failure Mode 4: Vendor selection misalignment

This is the most avoidable failure mode. Firms select AI vendors based on demo impressions rather than integration capability. The demo looks impressive — the AI generates a beautiful IC memo in thirty seconds. Nobody asks what happens when the model encounters a ground lease with a non-standard reset mechanism. Nobody asks whether the output integrates with the firm's existing underwriting templates. Nobody asks about the verification workflow downstream.

> *The environment breaks before the tool does. When your underwriting model lives in one analyst's head, when your LP reporting requires three people and a Slack thread to execute, no AI system on earth can operate effectively in that environment.*

The Verification Tax

There is a specific cost that compounds when AI is deployed without adequate infrastructure. I call it the Verification Tax: the hidden cost of poorly implemented automation that creates more work than it saves.

The pattern is consistent: someone discovers that AI can draft memos or generate summaries. The output looks impressive at first glance. Then a senior person reviews it and starts finding problems. The comps are from the wrong submarket. The cap rate does not match current conditions. The market narrative is generic. Now the senior person is spending time correcting AI output instead of reviewing human work product. The total time invested — the analyst generating the draft plus the senior person fixing it — is often more than if the analyst had done it the traditional way.

The fix is not more AI. It is better AI infrastructure. Start with workflows where the cost of error is low and the time savings are high. Deal screening. Report assembly. Portfolio monitoring. Build deeper as the system earns trust through demonstrated accuracy.

Understanding failure modes is necessary but not sufficient. The next section provides the implementation sequence — what to do first, in what order, and what to expect at each stage.

10

The 90-Day Implementation Roadmap

What to do first, in what order, and what to expect at each stage — from diagnostic through first production deployments.

Month 1: Diagnostic and data audit

The first month is about seeing clearly. No technology purchases. No vendor demos. Just an honest assessment of where the firm stands across the ten AIM Index dimensions, and a clear picture of what the Coordination Tax is costing.

Key activities

Score each of the ten AIM Index dimensions through structured interviews with the CIO, investment lead, asset management lead, IR lead, and operations lead. Map the current technology stack: every system, what data it holds, and the authoritative source for each data type. Estimate the Coordination Tax using the firm's own data — total G&A, team hours, and the coordination-to-judgment split. Identify the binding constraint: the lowest-scoring AIM dimension that limits everything else.

Month 1 deliverable

A scored assessment, a coordination tax estimate, and a prioritized remediation roadmap. This deliverable has standalone value regardless of whether the firm proceeds with implementation.

Month 2: First two workflow deployments

The second month targets the two highest-leverage workflows identified in the diagnostic — typically deal screening and report assembly. These are chosen because they have the highest ratio of coordination time to judgment time, the lowest cost of error, and the most immediate time savings.

Key activities

Build the Data Spine Layer 1 for the selected workflows: standardize the data definitions, formats, and field names. Deploy the first automated workflow with a human-review gate. Measure the actual time savings against the diagnostic baseline. Document what works and what does not — honestly.

What to expect

The first deployments will be imperfect. The value is not in the perfection of the output but in the learning: where does the data break? Where does the process need human judgment? Where is the AI output reliable enough to trust? These answers shape everything that follows.

Month 3: Infrastructure buildout and team alignment

The third month shifts from individual workflow automation to infrastructure. This is where the Data Spine layers get built, the governance framework gets formalized, and the team begins operating with the new processes as standard.

Key activities

Extend the Data Spine to cover portfolio-wide data flows. Formalize the AI governance policy (Green/Yellow/Red zones from Section 7). Train the team on the new workflows and the trust framework. Begin the second pair of workflow deployments. Establish success metrics and a measurement cadence.

Success metrics at 90 days

Metric	Baseline	90-Day Target
Deal screening time per opportunity	2–4 hours	30–60 minutes
Quarterly LP report assembly	3–4 weeks	1–2 weeks
Data reconciliation discrepancies	Unknown (not tracked)	Identified and quantified
AIM Index score (binding constraint)	Level 1–2	Level 2–3
AI governance policy	Does not exist	Documented and shared with team

The 90-day roadmap is the starting point — not the destination. The final section examines what happens beyond the first quarter: why the window for building this advantage is narrowing, and what the competitive landscape looks like for firms that move now versus those that wait.

> *Performance attracts capital. Capital enables acquisitions. Acquisitions build track record and train the system. The system improves future performance. The firms that close this loop tightest are the ones that scale fastest and build a durable competitive advantage — a true moat.*

11

Future Outlook and Competitive Implications

The bifurcation that is already underway — and why the window for building institutional infrastructure is narrowing.

The K-shaped sorting

The private equity industry is sorting itself. Bain's 2026 report describes a K-shaped recovery: a subset of elite funds on the upswing while everyone else struggles. The sorting mechanism is not deal quality or market timing. It is operational infrastructure.

The funds raising capital fastest have both strong returns and strong distributions. They can articulate a truly defensible, truly differentiated strategy — and prove with data how well the firm can execute.[1] The funds struggling have weak DPI, no systematic edge, and no data-backed story for LPs. After the last major stress period, fund series with bottom-quartile DPI performance failed to reraise at dramatically higher rates than those with top-quartile distributions.[27]

> **6%**
>
> of organizations qualify as AI "high performers" — where AI contributes meaningfully to EBIT. These firms are 2x more likely to have redesigned workflows before selecting AI tools.[24]

The LP diligence bar is rising

ILPA's DDQ has expanded from 8 sections to 21 since its original release. The average PE fund now responds to more than 150 DDQs annually during fundraising — up 40% from five years ago.[7] Technology governance, cybersecurity posture, and operational infrastructure are no longer peripheral questions. They are central to the LP's evaluation of whether a GP can execute consistently across market cycles.

ILPA is developing supplementary DDQ modules specifically for real estate — adding questions about technology infrastructure, data governance, and operational systems that were previously optional.[30] The firms that have already built this infrastructure will answer these questions from a position of documented strength. The firms that have not will be answering due diligence in real time rather than from a prepared position.

The compounding advantage

The most powerful feature of a well-built AI platform is not speed. It is that it gets better with every transaction. Every deal memo, every IC discussion, every market call feeds back into the system. Every underwriting model refines the assumption library. Every quarterly report improves the templates. Every portfolio variance analysis sharpens the monitoring thresholds.

This is the Compounding Loop. It runs through five stages: cleaner data reduces coordination overhead; reduced overhead frees capacity for investment quality and LP relationships; better investment quality leads to stronger performance and easier fundraising; easier fundraising provides capital and confidence to invest further in infrastructure; and stronger infrastructure enables more sophisticated deployment, which further reduces coordination overhead.

No competitor can buy this off the shelf. The combination of your criteria, your data, your history, and your judgment encoded into operational infrastructure is unique to your firm. That is the moat — and it deepens with every quarter you operate.

> *The firms building infrastructure now will not be talking about it later. They will just be the ones getting the allocations.*

The window

The CRE PE industry moved from fewer than 5% of firms piloting AI in July 2023 to 88% by October 2025.[15] PropTech venture investment reached $16.7 billion in 2025, up 68% year-over-year, with AI-native companies growing at 42% annually.[22] The technology is accelerating. The question is not whether the tools exist. It is whether your firm's operational foundation is ready to absorb them.

The cost of inaction is not standing still. It is falling behind firms that are compounding their infrastructure advantage every quarter. The gap between leading firms and lagging firms does not close over time. It widens. The same dynamic that makes the Compounding Loop powerful for firms that build early makes it punishing for firms that wait.

Glossary

Term	Definition
AIM Index	AI Maturity Index. A proprietary 10-dimension diagnostic framework for evaluating a CRE PE firm's operational readiness for AI deployment.
AUM	Assets Under Management. Total market value of assets a firm manages on behalf of investors.
Compounding Loop	The self-reinforcing cycle: better data → reduced coordination → freed capacity → better performance → easier fundraising → more infrastructure investment.
Coordination Tax	The aggregate time, effort, and opportunity cost consumed by finding data, reconciling formats, and reassembling information — the cost of operating without a platform.
Data Spine	The four-layer data infrastructure model: master data model → automated connectors → validation rules → monitoring and AI application layer.
DDQ	Due Diligence Questionnaire. Standardized by ILPA (version 2.0, 21 sections) for LP evaluation of GP operations.
DPI	Distributions to Paid-In. Measures actual cash returned to LPs relative to their invested capital.

Term	Definition
Firm Intelligence	AI systems calibrated to a specific firm's criteria, data, and investment philosophy — Level 2 of the Three Levels of AI Integration.
GP / LP	General Partner / Limited Partner. The manager (GP) and the investors (LPs) in a private equity fund structure.
IC	Investment Committee. The decision-making body that approves acquisitions, dispositions, and other investment actions.
ILPA	Institutional Limited Partners Association. Sets standards for LP-GP transparency, governance, and diligence.
Last Twenty Percent	The human judgment that AI cannot replace: relationship navigation, ethical reasoning, conviction under uncertainty, cultural reading.
MOIC	Multiple on Invested Capital. Total value (distributions + remaining value) divided by total invested capital.
NAV	Net Asset Value. The estimated fair market value of a fund's portfolio, net of liabilities.
NOI	Net Operating Income. Gross rental income minus operating expenses, before debt service and capital expenditures.
PERE	Private Equity Real Estate. The institutional segment of real estate investment through fund structures.
Platform vs. Tool	The distinction between using AI as an ad hoc tool (Level 1) and building AI into the investment process as connected infrastructure (Level 3).

Term	Definition
RVPI	Residual Value to Paid-In. The remaining unrealized value in the portfolio relative to invested capital.
TVPI	Total Value to Paid-In. DPI + RVPI. The most comprehensive return metric for a PE fund.
Verification Tax	The hidden cost of poorly implemented AI that creates more review work than it saves — the penalty for deploying at Level 1 without Levels 2 and 3.

Chirag Hathiramani

Real Estate Consulting

Two decades on the principal side. $3 billion in transactions. The lesson was always the same: the infrastructure is the moat. Now I help the next generation of CRE PE firms build it — and compound an advantage no competitor can replicate.

chi@chiraghathiramani.com
540-220-8102
chiraghathiramani.com
linkedin.com/in/addrealvalue

SCHEDULE A DIAGNOSTIC CALL

chiraghathiramani.com