WORKWISE SOLUTIONS

WHITE PAPER

Build vs. Buy vs. Partner

A Decision Framework for AI Capability Acquisition in Mid-Market Private Equity

Dr. Leigh Coney

Founder & Principal Consultant, WorkWise Solutions

Q1 2026

workwisesolutions.org

Contact@workwisesolutions.org

Abstract

Mid-market private equity firms face a three-way decision when acquiring AI capabilities: build an internal team, buy off-the-shelf tools, or partner with a domain specialist. Each path carries cost structures, timelines, and risk profiles that vendor marketing systematically obscures. A firm that commits to building in-house before understanding the talent market may spend eighteen months and $1.2 million before producing a single production system. A firm that buys generic tools may achieve operational efficiency without competitive differentiation. A firm that partners without a plan for knowledge transfer creates an external dependency that compounds over time.

This paper presents the Build-Buy-Partner Decision Matrix (BBPDM), a quantified framework that evaluates four variables - AI demand continuity, data proprietary value, talent competitiveness, and time-to-value pressure - to determine which acquisition path, or combination of paths, is optimal for a given firm. Drawing on compensation data, implementation timelines, and cost benchmarks from the current market, the paper provides specific guidance for general partners, operating partners, and chief technology officers at mid-market PE firms ($200M-$2B AUM). The analysis concludes that a phased hybrid model - partnering for speed and expertise, hiring for ownership and continuity, buying for commoditized functions - typically outperforms any pure path for firms in this segment.

Keywords: AI strategy, build vs. buy, private equity, AI talent, mid-market PE, AI implementation, technology acquisition, PE operating model, AI partner, hybrid model, decision framework, AI capability, PE technology strategy

Table of Contents

1. Introduction: The Question Every GP Is Asking

Every private equity firm evaluating AI eventually arrives at the same question: should we build this ourselves or buy it from a vendor? The question is understandable. It is also incomplete. There are three paths to AI capability, not two, and the optimal choice depends on variables that most firms have not quantified. The result is that the decision gets made by default - whoever has the most compelling pitch in a partner meeting, whichever vendor lands the best demo slot at the annual conference, or whichever operating partner has the strongest opinion about in-house technology teams.

This paper is not a sales pitch for any path. Building in-house is the right answer for some firms. Buying off-the-shelf tools is the right answer for others. Partnering with a specialist is the right answer for many. And for the majority of mid-market PE firms - those managing between $200 million and $2 billion in assets under management - a structured hybrid of all three paths typically outperforms any single approach. But the right hybrid depends on four specific variables that this paper will quantify.

The urgency of getting this decision right has increased substantially. Jensen Huang, NVIDIA's CEO, has described AI as essential infrastructure comparable to electricity or roads, arguing that every organization should be building AI capabilities into its operating model now rather than waiting for perfect conditions (Huang, 2026). The Stanford HAI AI Index Report 2025 documented that 78% of organizations reported using AI in at least one business function, up from 55% the prior year, with total corporate AI investment reaching $252.3 billion in 2024 (Stanford HAI, 2025). For PE firms specifically, the Bain & Company and StepStone Group 2026 GP Outlook found that roughly 40% of GPs do not expect material AI financial impact in 2026 - a finding that suggests not a lack of AI interest but a lack of clarity about how to acquire AI capabilities effectively.

This paper provides that clarity. The thesis is direct: for mid-market PE firms, a structured hybrid model typically outperforms pure build or pure buy. But the right hybrid depends on four variables - AI demand continuity, data proprietary value, talent competitiveness, and time-to-value pressure - that the Build-Buy-Partner Decision Matrix (BBPDM) is designed to quantify.

2. The Three Paths: An Honest Assessment

Before introducing the decision framework, it is necessary to assess each path on its actual merits and costs - not as vendors or advocates present them, but as firms experience them in practice. Each path has legitimate strengths, genuine weaknesses, and cost structures that are frequently misrepresented.

2.1 Path 1: Build In-House

Building an internal AI team means hiring the engineers, data scientists, and infrastructure specialists who will design, build, deploy, and maintain AI systems within the firm. It offers the greatest degree of control and, over time, the deepest institutional integration. It is also the most expensive path, the slowest to produce results, and the most vulnerable to a single structural risk: talent retention.

Minimum Viable Team and Costs

A production-capable AI team for a PE firm requires, at minimum, three to five people. The core roles and their current market compensation ranges (base salary, United States, 2025-2026 data) are as follows:

Role	Base Salary Range	Primary Function
ML/AI Engineer	$180,000 - $350,000	Model development, fine-tuning, prompt engineering, system architecture
Data Engineer	$150,000 - $280,000	Data pipelines, integration, quality assurance, infrastructure
Domain Translator	$160,000 - $300,000	Bridges AI capabilities with PE workflows, defines requirements, validates outputs
DevOps/MLOps	$140,000 - $250,000	Deployment, monitoring, scaling, security, compliance

Fully loaded cost - including benefits, equity, recruiting fees, infrastructure (cloud compute, API costs, development tools), and management overhead - adds 20-25% to base compensation. This places the annual cost of a minimum viable AI team at $1.2 million to $2.4 million per year, depending on seniority and location.

The ramp-up timeline is equally significant. Even after hiring is complete - which itself typically takes three to six months for specialized AI roles - the team requires six to twelve months before delivering its first production system. This timeline includes understanding the firm's data landscape, building integration pipelines, developing and testing models against real deal data, and deploying with the security and compliance requirements that PE workflows demand. A firm that begins hiring today can reasonably expect its first production AI capability in nine to eighteen months.

The Talent Retention Problem

The structural vulnerability of the build path is talent turnover. Aon's compensation and turnover studies covering 2022-2025 documented overall turnover in financial services declining from 18.5% to 15.5%, but noted that this aggregate masks significant variation by role type, with technology and AI specialists experiencing higher voluntary turnover driven by intense cross-industry competition for their skills (Aon, 2025). The Rise AI Talent Salary Report 2026 documented a 28% salary premium for AI roles over traditional technology positions, with AI-related job postings growing 257% since 2015 (Rise, 2026). When an AI engineer departs - and in a team of four, one departure every eighteen to twenty-four months is a reasonable planning assumption - the firm loses not just a headcount but the institutional knowledge embedded in the systems that person built. Rebuilding that knowledge takes months, during which system maintenance and development stall.

Strengths: Full control over architecture and priorities. Deep integration with proprietary data and workflows. Proprietary intellectual property. No external dependencies.

Weaknesses: Highest cost. Slowest time to value. Structural talent retention risk. Requires management capacity that most mid-market PE firms do not have available.

2.2 Path 2: Buy Off-the-Shelf

Buying means licensing commercially available AI-powered tools: deal sourcing platforms such as Grata, SourceScrub, and PitchBook; portfolio monitoring systems such as eFront and Allvue; reporting and investor relations tools such as Juniper Square. These tools are proven, immediately available, and maintained by their vendors.

Cost and Timeline

Annual licensing costs range from $50,000 to $500,000 depending on the scope of tools deployed, the number of users, and the level of data integration. Time to basic functionality is measured in days to weeks for standard features - a meaningful advantage over the build path. Implementation of more sophisticated features with firm-specific data integration may take two to four months.

The Differentiation Problem

The fundamental limitation of the buy path is competitive parity. Off-the-shelf tools are available to every firm willing to pay the licensing fee. If your firm and three competitors all use the same deal sourcing platform with the same underlying data, the platform provides operational efficiency but not competitive edge. The deals it surfaces for you are the same deals it surfaces for everyone else.

This matters in PE because the margin of advantage in deal sourcing is narrow. The firms that consistently access proprietary deal flow do so through relationships, thesis-driven origination, and analytical frameworks that reflect their specific investment strategy - none of which a generic tool can replicate. An off-the-shelf tool can help a team process information faster, but it cannot make the team’s analytical framework more distinctive.

A secondary concern is data security. Many commercial AI tools process firm data through shared vendor infrastructure. For PE firms handling confidential deal information under NDA, this creates a risk that must be evaluated against the vendor’s data handling practices, contractual protections, and technical architecture.

Strengths: Immediate availability. Proven at scale. Vendor maintains infrastructure, updates, and support. Lowest cost of the three paths.

Weaknesses: No competitive differentiation. Limited customization to firm-specific thesis and workflows. Data security concerns with shared infrastructure. Dependency on vendor roadmap for feature development.

2.3 Path 3: Partner with a Specialist

Partnering means engaging a firm with domain expertise in both AI engineering and PE workflows to design, build, and deploy custom AI systems. Unlike buying, the resulting systems are built to the firm's specific requirements. Unlike building in-house, the expertise is available immediately rather than after a multi-month hiring and ramp-up process.

Cost Structure

Partnership engagements for PE-focused AI systems typically follow a structured cost model. A discovery sprint - a two-to-four-week scoping engagement that maps the firm's data landscape, identifies the highest-ROI AI systems, and produces a technical roadmap - runs $25,000 to $75,000. A full build engagement for two to three production systems, delivered over eight to sixteen weeks, runs $150,000 to $500,000. Post-launch operational support on a monthly retainer runs $15,000 to $50,000 depending on the scope of ongoing capability development and maintenance.

The total cost of partnership for the first year - discovery, build, and retainer - typically falls between $250,000 and $700,000. This is roughly one-quarter to one-half the cost of a minimum viable in-house team, delivered in one-quarter of the time.

Strengths: PE domain expertise from day one. Custom to the firm's thesis and workflows. Production systems in weeks rather than months. Access to cross-client learning about what works in PE AI deployment.

Weaknesses: Creates external dependency. Institutional knowledge resides outside the firm unless deliberate transfer mechanisms are established. Ongoing cost for ongoing capability. Less control over development priorities than an in-house team.

3. The Four-Variable Decision Framework

3.1 Framework Overview

The Build-Buy-Partner Decision Matrix (BBPDM) evaluates a firm's optimal AI acquisition path across four variables. Each variable is scored on a 1-5 scale. The composite score profile - not a single aggregate number, but the pattern across all four variables - determines which path or combination of paths is most appropriate.

The framework is designed to be completed by a firm's general partner and operating partner (or CTO, where one exists) in a single working session. It does not require external consultants or technical expertise to apply. Its value lies in forcing the four questions that most firms skip when making AI acquisition decisions based on vendor presentations and peer anecdotes.

3.2 Variable 1: AI Demand Continuity

The question: How continuous is your firm's need for new AI capabilities?

Score 5: The firm anticipates a continuous stream of new AI projects - quarterly or more frequently. New fund strategies, new data sources, new analytical requirements, and evolving portfolio needs create ongoing demand for AI development. This level of demand can justify a permanent internal AI function because the team will have continuous productive work.

Score 1: The firm needs two to three specific AI systems built (deal screening, portfolio monitoring, IC memo support), after which the requirement shifts to operational management and periodic updates rather than continuous new development.

Implication: A score of 4-5 favors building, because the volume of work justifies the fixed cost. A score of 1-2 favors partnering or buying, because the demand profile is project-shaped rather than function-shaped. The mid-market reality is that most firms in the $200M-$2B range need three to five systems built and then managed - a project portfolio, not a permanent function. Most mid-market firms score 2-3 on this variable.

3.3 Variable 2: Data Proprietary Value

The question: Does your firm's data create a defensible competitive advantage?

Score 5: The firm has fifteen or more years of deal history, proprietary scoring models built from that history, unique market data from operational expertise in specific sectors, or LP relationships that generate exclusive deal flow. This data, combined with an AI team, could produce analytical capabilities that competitors cannot replicate.

Score 1: The firm primarily uses publicly available market data, third-party research, and standard deal flow channels. The data assets are not materially different from those available to competitors.

Implication: A high score favors building, because proprietary data is a moat that becomes more valuable when exploited with custom AI. A low score favors buying, because generic tools are adequate when the underlying data is generic. The key diagnostic question is: "Would a competitor with access to our data and an AI team create meaningful advantage over us?" If the answer is no, the data does not justify the cost of a custom build.

3.4 Variable 3: Talent Competitiveness

The question: Can your firm realistically attract and retain AI engineering talent?

Score 5: The firm is located in a technology hub (or offers fully remote positions), offers compensation competitive with technology companies, and can articulate a compelling mission that appeals to AI engineers - such as working on novel problems with proprietary data at a scale that a large technology company cannot offer.

Score 1: The firm is competing with Google, OpenAI, Anthropic, and well-funded AI startups for the same engineers, from a non-technology-hub office, offering 60-70% of market compensation, with a mission statement ("we do leveraged buyouts") that does not resonate with AI talent.

Implication: This is the variable where mid-market PE firms must be most honest with themselves. Most score 1-2. AI engineers want to work at AI companies. The PwC 2025 Global AI Jobs Barometer documented that workers with AI skills command a 32% wage premium over peers in the same role without AI skills - up from 25% the prior year - and that demand for AI talent continues to outstrip supply across all industries (PwC, 2025). A firm that scores 1-2 on this variable will experience chronic recruitment difficulty and elevated turnover regardless of how compelling the technical work is. Building under these conditions is not a strategy; it is an aspiration that will produce recurring frustration and wasted recruiting spend.

3.5 Variable 4: Time-to-Value Pressure

The question: How quickly must AI capabilities be operational?

Score 5: The firm's strategic plan allows twelve to eighteen months for AI capability development. Competitive pressure exists but is not immediate. The fund cycle provides a natural runway for building capability before it must be deployed on live deals.

Score 1: Competitive pressure, LP expectations, or a specific near-term opportunity requires AI capabilities within the current quarter. A nine-to-twelve-month build timeline is not viable.

Implication: The speed differential between paths is substantial. An in-house build typically requires nine to twelve months before first production output. A partnership engagement can deliver production systems in six to twelve weeks - a four-to-six-times speed advantage. Off-the-shelf tools can provide basic functionality in days. A firm scoring 1-2 on this variable cannot afford the build path regardless of how it scores on other variables.

Variable	High Score (4-5) Favors	Low Score (1-2) Favors
AI Demand Continuity	Build (continuous work justifies fixed team)	Partner or Buy (project-shaped demand)
Data Proprietary Value	Build (exploit the moat with custom AI)	Buy (generic tools for generic data)
Talent Competitiveness	Build (can attract and retain engineers)	Partner (access expertise without competing for talent)
Time-to-Value Pressure	Build (runway allows ramp-up)	Partner or Buy (speed is critical)

4. The Hybrid Model: What Most Mid-Market Firms Should Do

4.1 Why Pure Paths Fail for Mid-Market

When mid-market PE firms score themselves honestly on the BBPDM, a consistent pattern emerges. AI demand continuity is moderate (2-3): the firm needs several systems, not a permanent development function. Data proprietary value varies (2-4): many firms have valuable deal history but have not structured it for AI consumption. Talent competitiveness is low (1-2): the firm cannot realistically compete with technology companies for AI engineers. Time-to-value pressure is high (1-2): competitive dynamics and LP expectations demand near-term capability.

This profile - moderate demand, variable data value, low talent competitiveness, high time pressure - is incompatible with any pure path. Pure build fails because the firm cannot hire fast enough, retain reliably, or justify $1.2 million in annual cost before any output. Pure buy fails because it provides efficiency without differentiation, and the firm's proprietary data and thesis-specific workflows cannot be served by generic tools. Pure partner fails because it creates an external dependency without building internal capability - a dependency that becomes more expensive and more entrenched over time.

The hybrid model captures the strengths of each path while mitigating their individual weaknesses.

4.2 The Phased Hybrid (Recommended Timeline)

The phased hybrid model follows a structured sequence designed to deliver value early while building sustainable internal capability over time:

Weeks 1-4: Partner-Led Discovery Sprint. A specialist partner maps the firm's data landscape, identifies the highest-ROI AI systems based on the firm's specific workflows and investment thesis, and builds a working prototype. The discovery sprint produces a technical roadmap with prioritized deliverables, estimated timelines, and clear success criteria. This is the phase where the firm learns what is possible with its data before committing to a larger investment.

Weeks 4-16: Partner-Led Build. The partner builds the first two to three production AI systems - typically some combination of deal screening, portfolio monitoring, and IC memo support - using architecture designed for eventual internal ownership. The build phase uses the firm's actual data, integrates with existing tools and workflows, and deploys under the security requirements appropriate to PE deal data.

Weeks 6-10: Internal Hire Joins the Build. The firm hires one internal AI/data person - ideally someone with financial services experience and sufficient technical capability to manage (not necessarily build from scratch) AI systems. This person joins during the build phase, learning the architecture from inside the project rather than inheriting a completed system with no context. The hire is intentionally timed to overlap with the partner's most active build period.

Months 4-8: Gradual Ownership Transfer. The partner shifts from building to advising as the internal hire takes increasing ownership of day-to-day system management. The partner documents architecture decisions, operational procedures, and maintenance requirements. Knowledge transfer is deliberate and structured, not assumed.

Month 6 Onward: Internal Management with Partner Retainer. The internal team manages day-to-day operations: monitoring system performance, handling routine updates, and managing user requests. The partner remains on retainer for new capability development, architecture decisions, and the complex engineering work that falls outside the internal team's capacity. The retainer ensures access to expertise without the full cost of employing it.

Selective Buy: Commoditized Functions. Off-the-shelf tools are deployed for functions where customization adds no value: market data feeds, basic CRM, standard financial data providers, and generic reporting. The buy path is used for infrastructure, not for differentiation.

4.3 Cost Comparison: Three-Year View

The financial case for the hybrid model is compelling on a three-year basis:

Path	Year 1	Year 2	Year 3	Three-Year Total
Pure Build	$1.2M - $2.4M	$1.2M - $2.4M	$1.2M - $2.4M	$3.6M - $7.2M
Pure Buy	$150K - $500K	$150K - $500K	$150K - $500K	$450K - $1.5M
Pure Partner	$400K - $700K	$180K - $600K	$180K - $600K	$760K - $1.9M
Hybrid Model	$500K - $800K	$400K - $700K	$350K - $650K	$1.25M - $2.15M

The hybrid model costs approximately one-third to one-half of pure build on a three-year basis. More importantly, it delivers its first production system in weeks eight to twelve, compared to months nine to twelve for pure build. The pure buy path is cheaper, but the cost comparison is misleading: it produces efficiency without differentiation, a distinction that matters when fund returns depend on analytical edge rather than operational speed.

The hybrid's cost trajectory also improves over time. As the internal hire assumes operational management, the partner retainer decreases. By year three, the firm has production AI systems, an internal person who understands them, and access to specialist expertise on a cost-effective retainer - a sustainable operating model rather than a recurring capital commitment.

4.4 Risk Mitigation

The hybrid model hedges against the specific risks that cause pure paths to fail. Talent turnover is mitigated because the partner retains architecture knowledge and can bridge operational gaps during internal transitions. Scope creep is controlled because the partner provides external discipline - a defined engagement scope, fixed deliverables, and structured milestones that prevent the open-ended development drift common in internal projects. Technology shifts are managed because the partner stays current across multiple client engagements, bringing cross-firm learning about which approaches work in production PE environments and which fail.

The most significant risk of the hybrid model is the transition itself. If the internal hire is not sufficiently capable, or if knowledge transfer is not deliberate and structured, the firm can find itself dependent on the partner for operational management rather than strategic guidance. The mitigation is to define the transfer milestones explicitly from the engagement's outset and to hire the internal person early enough in the build phase that they accumulate genuine understanding rather than receiving a post-hoc handover document.

5. When the Framework Says "Build"

For some firms, the BBPDM clearly favors building in-house. The profile is consistent: these are typically larger firms (generally $5 billion or more in AUM, though the relevant variable is demand continuity rather than fund size) with a genuine continuous stream of AI development needs across multiple funds and strategies.

These firms can pay $300,000 or more in base compensation for senior ML engineers and compete credibly with technology companies on total package. They have fifteen or more years of proprietary deal data, proprietary scoring models, or unique sector expertise that creates a defensible analytical advantage when combined with custom AI systems. Their strategic timeline allows nine to twelve months of ramp-up before the first production system is required. And critically, they have management bandwidth to lead an AI function - a partner or operating principal who will own the team's priorities, resolve conflicts between the AI team's technical ambitions and the deal team's practical needs, and advocate for the function's budget in annual planning.

Even in this scenario, the recommendation is to start with a partner engagement to establish the architecture. Hiring against an existing technical blueprint is faster, produces better role definitions, and reduces the risk that a newly hired team spends its first six months debating architecture rather than building systems. The partner defines the architecture; the internal team builds and extends it.

> *The test is simple: if you can sustain $1.2M-$2.4M in annual cost, tolerate nine to twelve months before first output, and realistically attract talent that Google and OpenAI also want, building is viable. For most mid-market firms, one or more of these conditions fails.*

6. When the Framework Says "Buy"

At the other end of the spectrum, some firms are best served by buying off-the-shelf tools and forgoing custom AI development entirely. The BBPDM profile is equally consistent: these firms have generic data needs served by standard market data, third-party research, and common deal flow channels. Their analytical requirements do not extend beyond what commercially available platforms provide. Their teams are small - typically fewer than fifteen people - without the operational capacity to manage custom AI systems even if someone else builds them. Their AI budget is below $200,000 per year.

For these firms, the buy path is correct. Commercial deal sourcing, portfolio monitoring, and reporting tools provide genuine operational efficiency improvements. The fact that competitors have access to the same tools is less material when the firm's competitive advantage derives from relationship networks, sector expertise, or operational value creation capabilities rather than from analytical differentiation in screening or diligence.

Two warnings accompany the buy recommendation. First, buying creates dependency on the vendor's product roadmap. If the vendor does not prioritize PE-specific features - and most horizontal AI platforms do not, because PE is a small addressable market relative to their total customer base - the firm waits for capabilities that may never arrive. Second, the buy path should be a deliberate choice, not a default. Firms that buy because they have not evaluated the alternatives are not making a strategic decision; they are avoiding one. The BBPDM exists to ensure the choice is informed.

7. Conclusion and Recommendations

The build-versus-buy framing is a false binary that leads mid-market PE firms toward suboptimal AI acquisition decisions. The question is not whether to build or buy. It is which combination of building, buying, and partnering produces the fastest path to AI capability without creating dependencies the firm cannot manage.

For most mid-market PE firms, the answer is a phased hybrid: partner for speed and domain expertise, hire for ownership and continuity, buy for commoditized functions where customization adds no value. The phased hybrid delivers production AI systems in weeks rather than months, at roughly one-third the cost of a pure build, while building sustainable internal capability over time.

The recommendations are as follows. First, score your firm on the four BBPDM variables. Be honest. Most mid-market firms overestimate their talent competitiveness and underestimate their time-to-value pressure. Second, if the profile points to hybrid - and for most mid-market firms it will - begin with a discovery sprint. The sprint costs less than a single month of an internal AI team, delivers a technical roadmap and working prototype, and produces the information needed to make informed decisions about the subsequent build phase. Third, time the internal hire to overlap with the partner build, not to follow it. The single most common failure mode in hybrid implementations is hiring after the build is complete, producing an internal person who inherits systems they do not understand. Fourth, use the buy path deliberately for commoditized functions, but do not confuse efficiency tools with competitive advantage.

> *The question is not whether to use AI. The question is which acquisition path gets you there fastest without creating dependencies you cannot manage.*

References

Aon. (2025). Salary Increase and Turnover Study 2022-2025. Aon Human Capital Solutions.

Bain & Company and StepStone Group. (2026). 2026 Private Equity GP Outlook.

Coney, L. (2026a). Measuring AI ROI in private equity: A framework for decision velocity vs. decision quality. WorkWise Solutions White Paper Series.

Coney, L. (2026b). Agentic AI in private equity: Multi-agent orchestration for end-to-end deal workflows. WorkWise Solutions White Paper Series.

Coney, L. (2026c). AI governance across the deal lifecycle: From sourcing through portfolio monitoring. WorkWise Solutions White Paper Series. DOI: 10.2139/ssrn.6274559

Huang, J. (2026). Remarks at the World Economic Forum Annual Meeting, Davos, Switzerland, January 2026. NVIDIA Corporation.

PwC. (2025). 2025 Global AI Jobs Barometer. PricewaterhouseCoopers International.

Rise. (2026). AI Talent Salary Report 2026. Rise Works.

Stanford Institute for Human-Centered Artificial Intelligence. (2025). The 2025 AI Index Report. Stanford University. https://hai.stanford.edu/ai-index/2025-ai-index-report

About the Author

Dr. Leigh Coney is the Founder and Principal Consultant of WorkWise Solutions. With a PhD in Organizational Psychology, Dr. Coney has spent over a decade at the intersection of AI, behavioral science, and organizational design. His research focuses on decision-making frameworks in high-stakes environments, with particular attention to why sophisticated AI systems fail to achieve adoption and how their impact on organizational performance can be rigorously measured.

This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include "Closing the Accountability Gap" (December 2025), "Combating Automation Complacency in Financial Due Diligence" (Q1 2026), "The Skill Erosion Paradox" (Q1 2026), "AI Governance Across the Deal Lifecycle" (Q1 2026), "Measuring AI ROI in Private Equity" (Q1 2026), and "Agentic AI in Private Equity" (Q1 2026).

About WorkWise Solutions

WorkWise Solutions builds secure, purpose-built AI systems for private equity, venture capital, and investment banking firms. The firm specializes in zero-retention AI architecture that ensures proprietary deal flow and portfolio data never train public models.

WorkWise’s approach is grounded in a core insight: most AI implementations fail not because of technology but because of broken workflows and poor adoption strategies. Every engagement integrates behavioral science and organizational psychology into the technical design, ensuring that AI systems become invisible, indispensable parts of how investment teams actually work.

The Discovery Sprint - a two-to-four-week scoping engagement that maps your data landscape, identifies your highest-ROI AI systems, and produces a working prototype - is the recommended first step regardless of which long-term acquisition path your firm ultimately chooses.

Contact:

Contact@workwisesolutions.org | workwisesolutions.org

Schedule a consultation: calendly.com/contact-atqf/30min