

WORKWISE SOLUTIONS

WHITE PAPER

AI for Roll-Ups

Standardizing Operations Across Acquired Companies

Dr. Leigh Coney

Founder & Principal Consultant, WorkWise Solutions

Q2 2026

workwisesolutions.org

Contact@workwisesolutions.org

Abstract

Roll-up and buy-and-build strategies have become the default operating model in private equity. Add-on acquisitions reached 76 percent of all sponsor-backed buyouts by the end of 2022 and have remained the majority of sponsor activity since, accounting for 58 percent of deal flow in 2025. The model is straightforward in theory: buy a platform company, acquire smaller targets at lower multiples, integrate them onto a common operating spine, and exit the combined entity at a higher multiple. The model is brutal in practice. KPMG's 2025 analysis of more than 3,000 public-to-public M&A deals found that 57.2 percent of acquirers ultimately destroyed shareholder value, and Harvard Business Review's durable finding that 70 to 90 percent of M&A deals fail to create shareholder value has held up across decades of research. Information technology and data fragmentation are repeatedly cited as core drivers of these integration failures.

This paper introduces the Roll-Up AI Standardization Framework (RUSF), a structured approach to deploying artificial intelligence specifically in roll-up settings. The framework is not about adding AI on top of a successful integration. It is about using AI to make the integration possible at all, and to compress the timeline from acquisition to standardized operations from years to quarters. The framework rests on three pillars: the integration spine (the data, finance, and master records that must be unified for the platform to function as one entity), the operating layer (the workflows that drive EBITDA, where AI delivers the bulk of measurable value), and the acquisition engine (the playbook itself, which becomes faster and more reliable with each deal as the operator captures and reuses what works).

The paper is intended for the operators and sponsors who run roll-up platforms: platform chief executives, chief operating officers, chief financial officers, chief information officers, integration leaders, and the operating partners and deal partners at the sponsor level who fund and oversee them. It addresses one central question: when the roll-up model depends on integration succeeding, and most integrations fail, how should AI be deployed to change those odds?

Keywords

AI for roll-ups, buy-and-build private equity, post-merger integration, add-on acquisitions, ERP standardization, data harmonization, master data management, platform integration, value creation, agentic AI portfolio company, integration playbook, operating partner, AI in M&A, KPMG Pulse of Private Equity, Bain buy-and-build

Table of Contents

1. Introduction: The Integration Problem at the Heart of Roll-Ups
2. The State of Roll-Ups in 2026
3. Why Roll-Ups Fail (and Why It Is Almost Always About Integration)
4. Why Roll-Ups Are an Unusually Strong Fit for AI
5. The Roll-Up AI Standardization Framework (RUSF)
6. Pillar I: The Integration Spine
7. Pillar II: The Operating Layer
8. Pillar III: The Acquisition Engine
9. Implementation: A Phased Roll-Up Operator Playbook
10. Common Failure Modes and How to Avoid Them
11. Conclusion and Recommendations

References

About the Author

About WorkWise Solutions

1. Introduction: The Integration Problem at the Heart of Roll-Ups

The roll-up has quietly become the way that most of private equity now operates. Sponsors no longer talk about whether to pursue add-on acquisitions; they talk about how many. Add-on acquisitions reached more than 76 percent of all private-equity-backed buyouts by the end of 2022, dipped only slightly in 2023, and have remained the majority of sponsor activity ever since (Goodwin, 2024).

Capstone Partners' Middle Market Private Equity Index found that add-on acquisitions accounted for 58 percent of sponsor activity in 2025, and the firm describes the buy-and-build model as having evolved from a distinct strategy to the industry's default operating framework (Capstone Partners, 2026).

KPMG's Pulse of Private Equity reports that buy-and-build strategies remained central across the EMA region in 2025, with add-on investment reaching a four-year high (KPMG, 2026).

The financial logic is clear. A platform acquired at twelve times EBITDA can grow by adding companies acquired at six or seven times EBITDA, with the combined entity ultimately valued at the platform multiple or higher. This multiple arbitrage is real, but it depends on a single condition: the acquired companies must actually be integrated into the platform. If they are not, what the sponsor owns is not a roll-up. It is a holding company, a collection of independent businesses that share a balance sheet and little else.

This is where the model breaks. Independent industry analysis estimates that more than two thirds of roll-up strategies fail to create any value for investors, with the underlying cause traced to integration failure rather than thesis failure (PMI Stack, 2026). The broader M&A picture is similarly stark. KPMG's 2025 study of more than 3,000 public-to-public M&A deals between 2012 and 2022 found that 57.2 percent of acquirers ultimately destroyed shareholder value, with total shareholder return dropping an average of 7.4 percent in the two years following close (KPMG, 2025). Harvard Business Review's widely cited finding that 70 to 90 percent of M&A deals fail to create shareholder value has held up across decades of research (HBR, 2011). Gartner research has found that 83 percent of data migrations fail or exceed their budgets, and that 84 percent of information technology integrations experience significant issues (Gartner data cited in PMI Stack, 2026). The pattern is consistent across studies: the deal closes, the synergy model is approved, and then the integration runs into reality.

Reality, in a roll-up, is unusually difficult. A platform with eight acquired companies is operating with up to eight charts of accounts, eight customer master files, eight enterprise resource planning systems, eight payroll providers, eight reporting formats, and eight ways of defining basic operational metrics. The work of harmonizing this is enormous. It is also the work most likely to be deprioritized when the next deal arrives. Each new acquisition adds a layer of complexity to a stack that was already incomplete.

This is the problem AI is genuinely well suited to address. Unlike many of the AI use cases in financial services, where adoption is a matter of incremental productivity gains, the AI use case in roll-ups is about doing work that is otherwise prohibitive at human scale. Reading and reconciling thousands of

supplier records across a dozen systems is not work that benefits from a 20 percent productivity uplift. It is work that either gets done or does not, and AI changes which side of that line the operator sits on.

This paper introduces the Roll-Up AI Standardization Framework (RUSF), a structured approach to deploying AI specifically in roll-up settings. It does not promise that AI makes integration easy. It argues that integration without AI has become impractical at the cadence at which today's most active sponsors are buying.

2. The State of Roll-Ups in 2026

Three patterns characterize the roll-up landscape entering 2026.

First, the model has consolidated. Add-on acquisitions are the majority of sponsor activity and have been for several years. Capstone Partners reports add-on share at 58 percent in 2025, down modestly from 61 percent in 2024 but still firmly the dominant deal type (Capstone Partners, 2026). KPMG's Q4 2025 data confirms that buy-and-build strategies remained key across the EMA region in 2025, with add-on investment hitting a four-year high (KPMG, 2026). The acquired company is now more often a small business being absorbed into a platform than a standalone target being held independently.

Second, the operating bar has risen sharply. Bain & Company's 2026 Global Private Equity Report describes the new era with the rule of thumb that twelve is the new five: today's deals demand the kind of EBITDA growth that used to define top-quartile performance (Bain & Company, 2026). With holding periods drifting toward seven years and distributions to limited partners holding below 15 percent of net asset value for four years running, sponsors cannot rely on multiple expansion or cheap leverage to drive returns. Operational improvement, including the operational improvement that comes from genuine post-merger integration, has become the primary path to returns rather than a supplement to them.

Third, the sectors most active in roll-ups are also the most fragmented and the most operationally complex. Healthcare, where physician practices, specialty clinics, and healthcare technology providers are being consolidated, is the leading roll-up sector (DealRoom, 2026). Consumer and residential services, including HVAC, plumbing, and adjacent trades, continued to be a major focus of private equity investment in 2025, with PitchBook data showing private equity firms and their portfolio companies accounting for the majority of transactional activity in HVAC M&A (Cherry Bekaert, 2026). Professional services, including accounting, wealth management, and consulting, have been a similarly active roll-up theatre. None of these sectors offers clean, modern, well-instrumented operations as a starting point. Most offer the opposite: small founder-led businesses with idiosyncratic systems, paper-based processes, and informal data.

The combination of these three patterns produces the central operational tension of the modern roll-up. The strategy demands the rapid integration of small, operationally messy targets, while the financial environment demands more demonstrable EBITDA improvement than ever, all under tighter holding periods. The operator who tries to do this without AI is competing against operators who are not.

3. Why Roll-Ups Fail (and Why It Is Almost Always About Integration)

Roll-up failure is not random. It follows a pattern that integration practitioners have documented for decades, and that several recent analyses have quantified specifically for the private equity context.

3.1 The Three Layers of Integration Debt

Integration debt accumulates in three layers, each more expensive to resolve than the last.

The first layer is data debt: inconsistent customer records, duplicated supplier records, incompatible product hierarchies, mismatched stock keeping units, and conflicting general ledger codes across the acquired companies. Each acquired company brings its own version of these records, and the platform has to choose: keep them separate, lift and shift them, or harmonize them. Lift and shift is the path most commonly chosen and the one most likely to fail. Industry analysis of post-merger ERP work emphasizes that disparate naming conventions, inconsistent hierarchies, and duplicated records disrupt reconciliation, delay reporting, and increase audit exposure across departments (KPC Team, 2026).

The second layer is process debt: different ways of quoting, invoicing, paying suppliers, recognizing revenue, closing the books, hiring, onboarding, and managing customer service. Process debt is harder to see than data debt but bites harder. A platform with eight acquired companies often has eight different definitions of when revenue is recognized.

The third layer is system debt: parallel ERPs, parallel CRMs, parallel payroll systems, parallel point-of-sale or job-management systems. System debt is the most visible, the most expensive to resolve, and the one most often promised to be fixed soon by an integration plan that never quite gets there.

3.2 The Time Mismatch

There is a chronic mismatch in roll-ups between how quickly synergies are promised and how quickly integration actually happens. Industry analysis finds that the first 100 days typically address only the most critical elements, including keeping email running, ensuring basic systems connectivity, and maintaining customer-facing operations, while full integration with a unified ERP, consolidated CRM, and rationalized technology stack typically takes twelve to eighteen months at minimum, and two to four years for complex technology companies (PMI Stack, 2026). Industry estimates also link 50 to 60 percent of synergy capture initiatives directly to information technology execution (PMI Stack, 2026).

The financial model assumes synergies materialize earlier than the technology timeline can deliver them. The pressure created by this mismatch is what causes operators to take shortcuts that produce integration debt. Each shortcut is rational at the moment it is taken. The accumulation is what kills the deal.

3.3 The Compounding Problem of Serial Acquirers

In a single acquisition, integration debt is bounded. In a roll-up, it compounds. Each new add-on adds its own data, processes, and systems to an integration that was already incomplete. Without deliberate counter-action, the platform's debt grows faster than its operating capacity to address it. This is the structural reason that more than two thirds of roll-ups fail to create value for investors (PMI Stack, 2026): not because the strategy is wrong, but because the operating discipline required to execute it cleanly across many acquisitions is rarely present.

AI does not eliminate this problem. It changes the cost structure of solving it. The next sections explain how.

4. Why Roll-Ups Are an Unusually Strong Fit for AI

Most enterprise AI use cases are about productivity at the margin. The roll-up use case is structurally different. Three characteristics of the roll-up environment make it an unusually strong fit for AI deployment, and each of them maps to a specific category of integration work.

4.1 Repetition Across Acquired Entities

Every roll-up involves repeating the same kinds of work across many acquired entities: cleaning up customer records, mapping general ledger accounts, normalizing supplier master data, consolidating product catalogs, harmonizing employee records, and re-coding historical financials onto the platform's chart of accounts. This is exactly the kind of repetitive structured work that modern AI handles well, and exactly the kind of work that scales poorly with human labour.

A platform integrating its eighth add-on is doing a more sophisticated version of work it did for the seventh, the sixth, and the fifth. An AI system can be trained on the seven prior integrations and treat the eighth as a known pattern. Human teams cannot retain that learning curve as cleanly because each integration project tends to be staffed differently.

4.2 High Volume, Low Margin of Error

Integration work is voluminous (tens of thousands of records to reconcile in a typical mid-market platform with several add-ons) but has very little tolerance for error. A duplicate customer record may not be visible during the integration, but it surfaces during the audit, the year-end close, the buyer due diligence, or the regulatory examination. AI is uniquely good at high-volume, low-error-tolerance pattern matching, particularly when paired with structured human review at the points where confidence is low.

4.3 The Standardization Premium

Standardization itself produces enterprise value at exit. Buyers and underwriters in 2026 expect clean audit trails, consistent historical reporting, and documented controls and governance, and the absence of these things directly affects valuation and timing (DFIN, 2026). Industry guidance now emphasizes embedding SEC-ready reporting, audit-guide documentation, and governance controls immediately after acquisition rather than shortly before exit. AI deployments that compress the standardization timeline therefore do more than reduce internal cost; they protect and often increase the exit multiple.

Combined, these three characteristics produce a use case profile that differs sharply from the typical enterprise AI deployment. This matters because it changes the economics of the investment. Where most AI projects need to justify themselves through productivity gains, AI in roll-ups frequently justifies itself through making otherwise-impractical work possible at all.

5. The Roll-Up AI Standardization Framework (RUSF)

The RUSF is a three pillar framework for the operators and sponsors running roll-up platforms. It rests on a single premise: in a roll-up, the operating model is integration, and integration is a data and standardization problem before it is anything else. AI investment should follow that priority.

Pillar	What It Standardizes	Where AI Earns Its Place
I. The Integration Spine	Master data, chart of accounts, financial close, basic reporting. The systems and records that must be unified for the platform to function as one company.	Master data harmonization, chart of accounts mapping, automated record reconciliation, document classification, anomaly detection in legacy financials.
II. The Operating Layer	Customer-facing and revenue-driving processes: pricing, quoting, scheduling, customer service, sales enablement, procurement, and operational analytics.	Pricing optimization, quote generation, customer service automation, demand forecasting, churn detection, procurement intelligence, cross-sell identification.
III. The Acquisition Engine	The acquisition and integration playbook itself: target screening, diligence, day-one preparation, integration execution, and the institutional learning that compounds across deals.	Target sourcing, automated diligence, integration playbook execution, transferable templates, retained learning from prior deals.

Pillar I (The Integration Spine) addresses the foundation: the data, master records, chart of accounts, and financial close that must be unified for the platform to operate as a single company. AI investment in this pillar is unglamorous but indispensable. Without a working spine, the platform is not a roll-up.

Pillar II (The Operating Layer) addresses the workflows that drive EBITDA: pricing, customer service, demand forecasting, procurement, and the operational analytics that turn integrated data into operating decisions. This is where the bulk of measurable value creation lives, but only after the spine is functional.

Pillar III (The Acquisition Engine) addresses the playbook itself. A roll-up that integrates one company well has done good work. A roll-up that builds a repeatable engine for integrating any future acquisition has built a strategic asset. AI is what makes the engine learn.

The framework consists of twelve design principles distributed across the three pillars. They are intended to be applied in sequence, because pillar two depends on pillar one and pillar three depends on both.

6. Pillar I: The Integration Spine

The integration spine is the work that has to be done for the platform to operate as a single company. Most of it is invisible to customers and to limited partners, but its absence is what causes integration failure. The principles below describe how AI should be deployed in this layer.

Principle 1: Treat Master Data as the First Investment

Customer records, supplier records, employee records, product or service hierarchies, and the chart of accounts are the platform's master data. Every downstream system depends on them. In a roll-up, master data from each acquired company arrives in a different shape and quality. The first AI investment should be in master data management: automated deduplication, fuzzy matching across naming conventions, hierarchy alignment, and the establishment of golden records that all systems eventually reference.

This is unglamorous work and does not feature in any pitch deck. It is also the single highest-leverage AI investment in a roll-up. Industry post-merger integration analysis is consistent: the most common cause of integration failure is data fragmentation that was not addressed in time (KPC Team, 2026; PMI Stack, 2026).

Principle 2: Automate Chart of Accounts Mapping

Every acquired company has its own chart of accounts. The platform needs one. Manually mapping a target's chart to the platform's chart, account by account, is one of the most time-consuming integration tasks and one of the most error-prone. AI systems can ingest the target's chart, propose mappings to the platform's chart, flag accounts that do not have an obvious match, and generate the historical restatement of the target's financials onto the platform's chart. This compresses what is often a multi-month exercise into days of preparation followed by structured human review.

Critically, the AI does not replace the human review. It changes what the human reviews from the entire chart to the small subset of accounts where the model is uncertain. This is the right division of labour for this kind of work.

Principle 3: Make Financial Close a Continuous Process

Most roll-ups close the books on a monthly cadence that lengthens as more entities are added. AI agents change this dynamic. Continuous reconciliation agents can match invoices, payments, and journal entries as they arrive, surface variances within hours rather than weeks, and prepare the close as a steady-state output rather than a periodic crunch. Industry analysis describes this as the move from reconciliation as event to reconciliation as background process. The platform reaches the point where adding a ninth acquired company does not lengthen the monthly close.

Principle 4: Standardize the Reporting Layer Before the Operating Layer

Before AI is deployed in customer-facing or operational workflows, the platform should be able to produce one consistent set of financial and operational reports across all acquired companies. This is not a deferred goal; it is the precondition for any operating-layer AI to be trustworthy. AI deployed on inconsistent reporting data produces inconsistent recommendations. Recent industry analysis describes this failure mode bluntly: AI can become a multiplier of operational weaknesses and risk exposure when implemented on inadequate or fragmented data infrastructure (PE Professional, 2026).

7. Pillar II: The Operating Layer

Once the integration spine is functional, AI investment moves to the operating layer: the workflows that drive EBITDA. The principles in this pillar are common across many private equity AI playbooks, but their application in a roll-up has a specific shape because the operator has multiple acquired companies to deploy across in parallel.

Principle 5: Standardize Pricing and Quoting Across Acquired Companies

In services-based roll-ups, including HVAC, plumbing, dental practices, and similar consumer and professional services, the acquired companies typically arrive with their own pricing logic, quote templates, and discount practices. This is one of the largest hidden EBITDA opportunities in a roll-up. AI-enabled pricing tools that ingest historical job data, regional benchmarks, and customer characteristics can produce consistent, optimized pricing across the platform without requiring every acquired company to abandon its local relationships.

A common pattern is to deploy the pricing tool first at one acquired company, validate the EBITDA impact, document the playbook, and roll out across the remainder of the platform on a structured cadence. This is consistent with the broader pattern in successful AI deployment, which is to apply AI to specific workflows supported by reliable data and driven by the business rather than by the technology team.

Principle 6: Concentrate Customer Service Through AI

Acquired companies usually arrive with their own customer service operations. Consolidating these into a single AI-augmented operation captures real cost savings and produces a more consistent customer experience. The right pattern is rarely to replace human agents wholesale; it is to deploy AI for triage, routine handling, and after-hours coverage, while reserving human attention for the cases where it matters. Service consolidation also produces a single dataset of customer issues across the platform, which becomes a high-value input to procurement, product, and operational decisions.

Principle 7: Use Demand Forecasting and Operational Analytics to Reallocate Across the Platform

A platform with eight acquired companies has eight regional or service-line operations. AI-enabled forecasting can identify capacity imbalances, surface cross-platform optimization opportunities (a slow week in one region, a backlog in another), and inform staffing and inventory decisions that no single acquired company could see on its own. This is value the platform produces that no acquired company could have produced as a standalone.

Principle 8: Procurement Consolidation Is the Most Underrated EBITDA Lever

Procurement is consistently undervalued in roll-up planning and consistently overdelivers in execution. AI agents can analyze the combined spend across all acquired companies, identify duplicate suppliers and pricing inconsistencies, surface negotiation opportunities, and continuously monitor for contract

compliance once consolidated agreements are in place. This work is repetitive, high-volume, and well suited to AI. It also produces measurable EBITDA impact in the same fiscal year, which is rare in roll-up integration work.

8. Pillar III: The Acquisition Engine

The first two pillars optimize the platform that exists. The third pillar is what separates a roll-up that integrates well from a roll-up that becomes a programmatic acquirer. The principles here describe how AI is used not to operate, but to acquire.

Principle 9: AI for Target Sourcing in Fragmented Sectors

Roll-ups operate in fragmented sectors with thousands of small targets. Manual sourcing of those targets is slow and biased toward whatever shows up in deal-broker email inboxes. AI agents can scan the broader universe of potential targets continuously, score them against the platform's criteria (geography, revenue, service mix, ownership profile), and surface high-fit opportunities before they reach a competitive process. This compounds: the more targets the platform has integrated, the more accurate the scoring criteria become.

Principle 10: Integration-Aware Diligence

In a roll-up, diligence is not just about whether the target is a good business. It is about how cleanly the target will integrate. AI-assisted diligence agents can compare the target's chart of accounts, customer master, supplier master, and core systems against the platform's standards, producing an integration risk assessment alongside the financial diligence. This changes the question from will this deal close to how expensive will this integration be, which is the question the roll-up operator actually needs answered.

Principle 11: A Living Integration Playbook

Every acquired company integrated produces a body of knowledge: which mappings worked, which conversion rules applied, which exceptions arose, which change-management approaches succeeded. In most platforms, this knowledge sits in the heads of the integration team and is partially lost when team members rotate. AI tools can capture and operationalize this knowledge as a living integration playbook, where the platform's experience from each prior deal informs the playbook for the next.

This is the single most important habit for compounding value in a roll-up. The platform that integrates its eighth add-on faster, more reliably, and at lower cost than its first is operating from a strategic asset that the platform that does not has not built.

Principle 12: A Failure Registry

Not every integration task succeeds the first time. AI implementations sometimes mismap accounts, mishandle exceptions, or surface false positives at high rates. The platform that captures these failures honestly, in a registry that is reviewed quarterly, learns faster than the platform that minimizes them in reporting. This principle parallels the failure registry recommendation in the Behavioural Governance Framework and the Portfolio AI Deployment Framework, and applies with particular force in a roll-up because the same failure pattern is otherwise repeated across each new acquisition.

9. Implementation: A Phased Roll-Up Operator Playbook

The RUSF is most useful when applied as a sequence rather than all at once. The following timeline is calibrated to a platform that has completed at least its first two or three add-ons and is committing to a serial acquisition strategy.

Phase 1: Stabilize the Spine (Months 1 to 4)

Deploy AI investment in the integration spine first. Build or procure master data management, automated chart of accounts mapping, and continuous reconciliation. Reach the point where the platform can produce one consistent set of financial reports across all entities, on a monthly cadence that does not lengthen with each new acquisition. This is the precondition for everything that follows.

Phase 2: Operating Layer Quick Wins (Months 4 to 9)

With the spine functional, deploy the operating-layer use cases that produce same-year EBITDA impact. Procurement consolidation and pricing standardization are usually the highest-leverage starting points. Customer service consolidation often follows. The pattern in each case is the same: deploy at one acquired company, validate the EBITDA impact, document the playbook, and roll out across the platform on a structured cadence rather than all at once.

Phase 3: Build the Acquisition Engine (Months 9 to 18)

Once the platform has demonstrated that it can integrate, operate, and improve an acquired company, the focus shifts to the engine. Deploy target sourcing AI. Build the integration-aware diligence capability. Operationalize the living integration playbook so that the platform's experience compounds. By the end of this phase, the platform should be measurably faster at integrating its next acquisition than it was at integrating its previous one.

Phase 4: Continuous Compounding (Ongoing)

From this point onward, every new acquisition feeds the engine. Master data harmonization gets faster because the AI has seen more variations. Pricing rolls out faster because the playbook is mature. The integration playbook deepens with each deal. This is the state in which the roll-up model produces the multiple-arbitrage returns it was designed to produce, because the operating discipline has caught up with the financial logic.

10. Common Failure Modes and How to Avoid Them

Several failure patterns appear frequently enough in roll-up AI deployments to warrant explicit naming.

Operating Before Integrating

The most common failure is deploying operating-layer AI before the integration spine is stable. The result is AI tools producing recommendations on the basis of inconsistent or unreliable data, which damages trust in the technology and in the integration team. The fix is sequencing: spine first, then operating layer, then engine.

Choosing the Wrong Integration Touch Level

Industry guidance distinguishes between high-touch (full ERP consolidation), medium-touch (unified financial spine with operationally flexible modules), and low-touch (lightweight integration with autonomous operations) approaches (PMI Stack, 2026). Operators often default to high-touch out of a desire for elegance, then discover that the timeline and disruption are not compatible with the deal cadence. Matching touch level to acquisition characteristics is one of the most important decisions in roll-up integration design.

Underestimating Change Management

AI tools that work technically still fail when the workforce will not use them. This is acutely true in roll-ups, where employees of acquired companies often experience the platform's standards as imposed rather than chosen. Industry analysis is consistent that successful integrations identify digital champions within acquired companies, deliver role-based training rather than generic demonstrations, and treat integration as a people problem with a technical component (PMI Stack, 2026). The same applies to AI deployments within those integrations.

Letting Each Add-On Reset the Playbook

Without deliberate capture, each acquired company's integration team works from scratch. The fix is the living integration playbook described in Principle 11. The platform that builds this once benefits across every subsequent deal.

Building When Buying Would Work

Roll-up platforms, particularly those with technology backgrounds, sometimes choose to build their own AI tooling for integration work. The MIT NANDA research is consistent across enterprise contexts that vendor and partnership solutions succeed roughly 67 percent of the time, while internal builds succeed only 33 percent of the time (MIT NANDA, 2025). For most roll-up operators, the buy-then-customize path is materially more likely to produce timely results than the build-from-scratch path.

11. Conclusion and Recommendations

The roll-up has become the dominant operating model in private equity, and most roll-ups underperform their integration plans. The problem is not that the strategy is wrong. It is that the operating discipline required to execute the strategy has not kept pace with the cadence at which today's most active sponsors are buying. Integration debt accumulates faster than integration capacity, and the result is value that exists in the financial model but not in the operating reality.

AI changes this only if it is deployed with discipline. The Roll-Up AI Standardization Framework provides a structured way to apply that discipline: stabilize the integration spine first, deploy operating-layer use cases for EBITDA impact second, and build the acquisition engine for compounding returns third. Each step depends on the one before it, and skipping any of them produces predictable failures.

The recommendations for sponsors and operators running roll-up platforms are as follows:

1. Sequence the spine before the operating layer. Master data, chart of accounts, and continuous close are unglamorous and indispensable. Without them, every operating-layer AI is unreliable.
2. Deploy procurement consolidation and pricing standardization first in the operating layer. Both produce same-year EBITDA impact and both have a strong AI fit.
3. Build a living integration playbook. The platform that integrates its eighth add-on faster than its first has a strategic asset. Without deliberate capture, that asset is never built.
4. Maintain a failure registry. The fastest way to compress the integration learning curve for the next deal is to be honest about what failed in the last one.
5. Choose buy over build. The data on success rates is decisive, and roll-up operators rarely have the time or the technology depth for build to be the right choice.
6. Match integration touch level to deal characteristics. High-touch is right sometimes. It is not right by default.
7. Treat AI deployment as part of the value creation plan, not an information technology project. Operating partners should fund and oversee it as they do any other value lever.
8. Position for exit by standardizing early. Buyers and underwriters will price the absence of clean reporting and governance into the exit multiple. AI-enabled standardization is one of the few investments that pays back both during the hold and at exit.

The roll-up that wins in 2026 and beyond is not the one that buys the most companies. **It is the one that integrates the companies it buys, faster and more reliably than the next operator. AI is what determines which side of that line a platform ends up on.**

References

Bain & Company. (2026). Welcome to a New Era: Global Private Equity Report 2026. Bain & Company.

Capstone Partners. (2026). Middle Market Private Equity Index Report and 2025 Middle Market M&A Valuations Index. Capstone Partners.

Cherry Bekaert. (2026). Private Equity Report: 2025 Trends and 2026 Outlook. Cherry Bekaert.

DealRoom. (2026). Private Equity Roll-Up Strategy: Complete Guide. DealRoom.

DFIN Solutions. (2026). Private Equity Trends 2026. Donnelley Financial Solutions.

Goodwin Procter. (2024). Use of Add-On Acquisitions in PE Is Likely to Continue. Goodwin Procter LLP.

Harvard Business Review. (2011). The Big Idea: The New M&A Playbook. C. Christensen, R. Alton, C. Rising, & A. Waldeck. Harvard Business Review.

KPC Team. (2026). Post Merger ERP System Integration: Protect Value After Acquisition. KPC Team.

KPMG International. (2025). All That Glitters Is Not Gold: M&A Synergies and Value in Public Acquisitions. KPMG International.

KPMG International. (2026). Q4 2025 Pulse of Private Equity. KPMG International.

MIT NANDA Initiative. (2025). The GenAI Divide: State of AI in Business 2025. Challapally, A., et al. MIT Media Lab Project NANDA.

PE Professional. (2026). Why Private Markets' AI Ambitions Are Hitting a Wall. PE Professional.

PMI Stack. (2026). 50+ Post-Merger Integration Statistics (2026); Post-Merger Integration Challenges; ERP Consolidation After Merger: The Roll-Up Operator's Field Guide. PMI Stack.

About the Author

Dr. Leigh Coney is the Founder and Principal Consultant of WorkWise Solutions. With a PhD in Organisational Psychology, Dr. Coney has spent over a decade at the intersection of AI, behavioural science, and organisational 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 governance and deployment systems can be designed to account for human cognitive and social dynamics.

Dr. Coney's work is distinguished by its integration of behavioural science with practical technology deployment. He advises private equity firms, family offices, private credit funds, independent sponsors, and investment banks on AI strategy, governance design, and psychology informed change management.

This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include Helping Portfolio Companies Deploy AI: A General Partner's Operating Partner Playbook (Q2 2026), Agentic AI Governance in Private Equity: A Behavioural Framework for Autonomous Decision Systems (Q1 2026), Measuring AI ROI in Private Equity: A Framework for Decision Velocity vs Decision Quality (Q1 2026), and The Skill Erosion Paradox: Preserving Analytical Capability in AI Augmented Teams (Q1 2026).

About WorkWise Solutions

WorkWise Solutions builds secure, purpose built AI systems for private equity, venture capital, family offices, private credit funds, and investment banking firms. The firm specialises 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 behavioural science and organisational psychology into the technical design, ensuring that AI systems become invisible, indispensable parts of how investment teams actually work.

Contact: Contact@workwisesolutions.org | workwisesolutions.org

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