AI Roll-Ups & Buyouts: An Analysis of Strategy, Execution, and Value Creation

Erik de Stefanis

Harvard Business School: Independent Project

Fall 2025

This report draws on conversations with more than 15 investors, operators, and founders. I've included quotes and observations from 11 of them below. Most of these conversations were off the record, which is the only reason people were willing to be as candid as they were about their own companies and the industries they're buying into.

To keep that confidence while still getting the value of their perspective onto the page, I've anonymized every quote and refer to each person by number. The table below gives a short description of who each interviewee is, so the reader has some context for where the opinions are coming from.

Overview of the interviewees:

Tag	Description
Interviewee 1	VC-turned-AI rollup investor
Interviewee 2	Ex-PE/VC investor, now AI-rollup investor
Interviewee 3	Founder at AI-rollup of accounting/tax service firms
Interviewee 4	Founder of AI-rollup
Interviewee 5	Ex-VC, LMM PE GP
Interviewee 6	Ex-VC, LMM PE GP
Interviewee 7	Searcher-turned-CEO of SMB
Interviewee 8	Ex-PE investor, now pursuing AI-buyouts
Interviewee 9	VC-turned-AI rollup investor
Interviewee 10	Ex-VC turned AI-buyout operator
Interviewee 11	VC-turned-AI rollup investor

Definition of AI Buyouts and Roll-Ups

The AI-enabled buyout and roll-up strategy is a hybrid. It borrows the control orientation of private equity and the technology-driven value creation of venture capital, and it applies both to the same asset. In practice, the model typically means acquiring traditional service businesses and transforming them through AI, with three explicit goals, i.e. expanding margins, capturing data advantages, and arbitraging the gap between the multiple a services business trades at and the multiple a software business trades at.

The distinction between a buyout and a roll-up is quite significant operationally. An buyout typically entails acquiring one primary, scaled asset. Subsequent add-ons are common, but not as key to the strategy as when pursuing a roll-up, where growth is predominatly inorganic, i.e. attained through subsequent acquisitions. This is also referred to as the “crown jewel” approach, where depth of implementation matters far more than breadth of consolidation.

A roll-up pursues a serial acquisition strategy, consolidating anywhere from five to a hundred-plus companies within a fragmented industry into a single AI-enhanced platform. When venture investors talk about this approach, they usually have a specific sequence in mind: build the product, pilot it with customers, demonstrate the automation and margin improvements, then identify buyout candidates, possibly among the pilots. That sequencing tells you something about how the buyout camp thinks.

	PE Model	Crown Jewel Buy-out	AI Roll-Up	VC Model
Importance of software engineering	Very low	**High**	**High**	Very high
Importance of financial engineering	Very high	**High (but episodic)**	**High**	Very low
Companies exposed to	1	**1–5**	**5–100**	>10,000
Exposure achieved through	Acquiring company	**Acquiring companies**	**Acquiring companies**	Selling products
Value capture per end company	Very high	**Very high**	**Very high**	Generally low

Roll-ups face a materially different set of execution problems. Integration complexity grows with every acquisition. Deploying talent across multiple sites becomes a real constraint, and consolidating data across systems that were never designed to talk to each other compounds the work. One operator put it well: buying multiple companies with heterogeneous operations forces you into one of two paths. Either you standardize workflows at the holding level, which only works

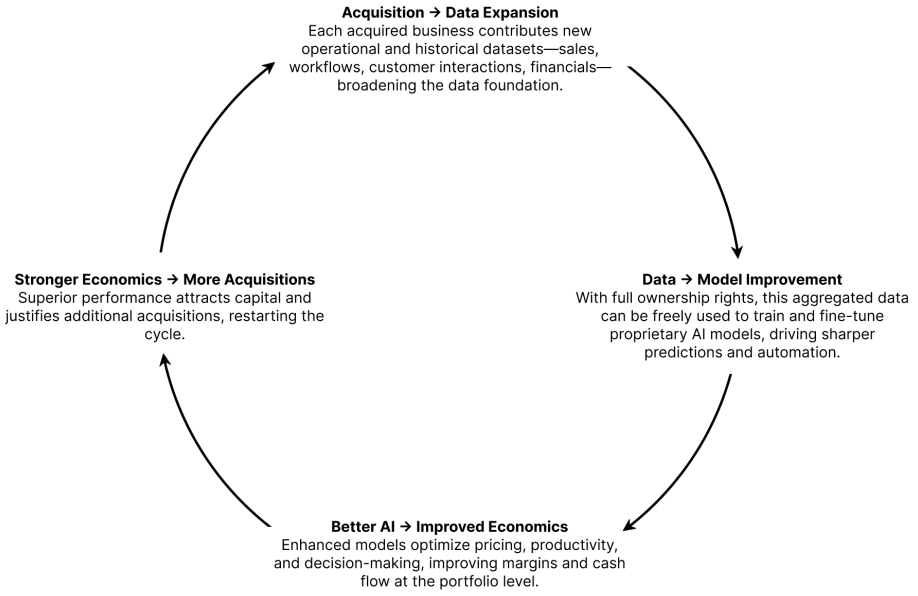
if the underlying operations are similar enough to begin with, or you customize target by target, which dramatically increases what you need from your product team.

The economic logic behind both approaches rests on a handful of interconnected value creation mechanisms.

Margin Expansion Through Labor Automation: The primary value driver is replacing expensive human labor with AI. Take a hypothetical accounting firm with 1,000 clients generating $600,000 in annual revenue at a 30% EBITDA margin. Through AI deployment, labor costs per client might fall from $30 to $18 per month while everything else stays flat. That one intervention takes EBITDA margins from 30% to 54%, which is an 80% increase in absolute cash flow.

	Example Accounting Firm	
	Pre-AI	**Post-AI**
Annual Revenue Per Client	$600	$600
Clients	1,000	1,000
Annual Revenue	**$600,000**	**$600,000**
Cost-to-Serve Per Client		
Labor Costs	$360	$216
Non-Labor Costs	$60	$60
Total Per Client	$420	$276
Total (Firm Level)	**$420,000**	**$276,000**
EBITDA	**$180,000**	**$324,000**
EBITDA Margin	*30%*	*54%*
Δ EBITDA		*$144,000*
Δ EBITDA Margin		*80%*
EBITDA Multiple	***4.5***	***10***
Enterprise Value	**$810,000**	**$3,240,000**

Data Access and Compounding Advantages: Ownership gives you direct access to operational data, which is the raw material for training better models. The theoretical flywheel runs as follows: more acquisitions generate richer datasets, richer datasets enable better AI, better AI improves unit economics, and better unit economics justify further acquisitions at increasingly favorable terms.



Interviewee 7 was blunt about how this plays out in reality: “Everyone talks about proprietary data as a moat. But honestly? Most SMBs have terrible data. It’s in spreadsheets, inconsistent formats, missing fields. You spend the first six months just getting clean data before you can even think about training models.” In my view, this undermines the data advantage thesis for most verticals, at least in the first couple of years.

Repeatable Transformation Playbooks: The model assumes that the AI implementation process can be standardized and repeated across acquisitions with a declining marginal cost each time. Interviewee 2 described a technical architecture designed explicitly for this: “There are basic functions that you can use repeatedly across all your platforms. Things like OCR, database infrastructure, vector databases, data pipelines, and agent orchestration frameworks.”

That composable infrastructure layer, built once, can then be deployed across every portfolio company at minimal incremental cost, with only the vertical-specific agents needing customization.

Multiple Arbitrage Between Services and Software: This is, to me, the most interesting part of the strategy. Traditional service businesses at the scale relevant to this report trade at 3-6x EBITDA. Pure software companies can command 10-100x. The thesis is that an AI-transformed service business should capture some portion of that premium and exit at 10-30x EBITDA instead of a services multiple.

Overcoming Implementation Resistance Through Ownership: This advantage gets underappreciated. Many service sectors have proven remarkably resistant to transformation, and the reason is rarely that better tools don't exist. The reason is cultural inertia, misaligned incentives, and the cost of change management.

Ownership removes the need to persuade an external customer to adopt your technology. As Interviewee 2 explained: “We’re only buying companies where the CEO has already bought into

the strategy and actually cares about the strategy. We're not replacing the CEO. We're making them roll 30 or 40% equity." Alignment through an ownership stake turns the implementation problem from one of persuasion into one of execution.

Venture Capital's Entry into the Strategy

Venture capital moving into buyouts and roll-ups is a real departure from how VC has traditionally invested. To evaluate whether the strategy is viable, it helps to understand why venture firms want to own services businesses instead of selling software to them, and how their fund economics differ from those of private equity.

Why VCs Are Pursuing Ownership Instead of Selling Software

Several factors converge here.

First, venture-backed software companies selling into SMB services markets have historically struggled with adoption barriers, long sales cycles, and limited willingness to pay. If you own the services business, the distribution problem disappears.

As Interviewee 1 put it: "AI enables 'selling work', bundled software and services. The target markets for services are a $16T market compared to a global global market of $1T for software." The addressable market for services is an order of magnitude larger than the one for software.

Interviewee 9 echoed this: "What became clear to us was where AI value creation will concentrate over the next 10 to 30 years in the services layer of the U.S. economy. After making three minority investments in AI-enabled services businesses in Fund I, all of which performed exceptionally well, it was only natural that Fund II would focus on this strategy."

Second, ownership lets the VC capture 100% of the AI-driven value creation. When you sell software, you get a licensing fee. When you own the business, you can capture the value created through margin expansion.

Third, venture investors have come to recognize that AI's impact is largest when it transforms an entire workflow, as opposed to when a single point solution gets adopted. That depth of transformation requires an operational integration that is very hard to achieve through arm's-length software sales.

Different Investment Philosophies

Traditional PE investors tend to approach AI-enabled strategies through the lens of operational excellence first and technology second. They understand the mechanics of acquiring businesses, integrating them, and improving margins. What they often lack is the technical sophistication to evaluate whether an AI implementation is feasible, or to execute one that is complex.

Within traditional PE firms, the cultural gap can be quite substantial. Most PE firms are used to hiring operators and consultants, less so engineers. Deal teams evaluate financials, customer

concentration, and contract structures. They do not evaluate API architectures, data pipelines, or model training requirements, which creates a blind spot when it comes time to judge whether a technical founder or CTO actually has the skills to build a sophisticated automation system.

Interviewee 6 described the typical PE posture: “The most common approach we will pursue is the traditional SMB rollup with an AI/tech enhancement layer. Focus on systematic optimization vs. relationship-based operations.” That is incremental optimization, and it is a long way from fundamental workflow transformation.

The operational complexity of all of this should not be underestimated. Interviewee 8, who moved from traditional PE into AI-enabled roll-ups, was not optimistic about the incumbents: “Middle market PE is fundamentally broken. Talent drains to mega funds or McKinsey. 70% of GP managers will disappear in 10 years as half are overleveraged and half are not innovating.” My read is that neither a pure VC background nor a pure PE background prepares someone for the hybrid demands of this strategy.

Core Assumptions Underpinning the AI Roll-Up Thesis

Every investment strategy rests on a set of foundational assumptions about where value comes from. For AI-enabled roll-ups, I count five core beliefs driving capital allocation. Going through each of them critically shows where the thesis holds up and where it runs into practical reality.

Margin Expansion Through Labor Automation

The central premise is that AI can systematically replace expensive knowledge workers with cheaper computational infrastructure. Interviewee 1 offered specific benchmarks after reviewing various industries and the share of labor hours spent on repeatable processes: “Call centers: 50–70% workflow automation. MSP/IT: ~38% automation. Bookkeeping: ~80% automation. Automation improves every 3–6 months.”

In a spreadsheet, the math is compelling, as it usually is. Replace a $150,000-per-year accountant with $30,000 in AI infrastructure and $120,000 in annual savings flows straight to EBITDA. Multiply that across a business with 50 accountants and you have $6 million in annual margin improvement. At a 10x EBITDA exit multiple, that single operational change creates $60 million in enterprise value.

		EBITDA Multiple					
		5x	10x	15x	20x	25x	30x
EBITDA Uplift	$1,000,000	$5,000,000	$10,000,000	$15,000,000	$20,000,000	$25,000,000	$30,000,000
	$2,000,000	$10,000,000	$20,000,000	$30,000,000	$40,000,000	$50,000,000	$60,000,000
	$3,000,000	$15,000,000	$30,000,000	$45,000,000	$60,000,000	$75,000,000	$90,000,000
	$4,000,000	$20,000,000	$40,000,000	$60,000,000	$80,000,000	$100,000,000	$120,000,000
	$5,000,000	$25,000,000	$50,000,000	$75,000,000	$100,000,000	$125,000,000	$150,000,000
	$6,000,000	$30,000,000	**$60,000,000**	$90,000,000	$120,000,000	$150,000,000	$180,000,000
	$7,000,000	$35,000,000	$70,000,000	$105,000,000	$140,000,000	$175,000,000	$210,000,000
	$8,000,000	$40,000,000	$80,000,000	$120,000,000	$160,000,000	$200,000,000	$240,000,000

In practice, the assumption is more nuanced. Not all white-collar labor is equally automatable. The operators I spoke with were consistent on this point: the white-collar work that LLMs currently excel at involves document review, data extraction and entry, simple knowledge retrieval, text-based communication, and materials creation. That said, model capabilities will continue to approve, and the scope of work at which they excel at will only increase. Work that requires physical presence, emotional intelligence, complex judgment under ambiguity, or high-stakes decisions with reputational risk attached remains hard to automate well.

Margin expansion also assumes that labor is a large enough share of the cost structure for automation to move EBITDA in a meaningful way. Interviewee 1 offered a sobering counterexample: “Medical scribe automation in dental practices sounds promising until one recognizes that scribes represent a small fraction of total practice costs, meaning even 100% automation yields minimal margin impact.”

Data Advantages and Model Improvement

The second core assumption is that direct ownership gives you access to proprietary operational data, which enables progressive model improvement and a defensible moat. The theoretical flywheel says each acquisition generates training data, the training data improves AI performance, better performance improves unit economics, and better unit economics fund the next acquisition with an even better model.

The data moat assumption runs into two substantial problems. The first is data quality, which fails along multiple dimensions at once: structural chaos, with customer data scattered across systems; format inconsistency, with the same data point recorded differently in different places; missing fields that were left blank because nobody needed them at the time; and quality degradation from data entered in a hurry and never cleaned up.

The implication for implementation timelines is significant. Pitch decks project margin improvement within three to six months of acquisition. The reality usually involves a “data infrastructure tax” of six to twelve months before any meaningful AI work can begin. That is a period of increased cash burn with zero AI-derived value creation, and it deepens the J-curve well beyond what investors often expect, which in turn creates timeline pressure that can erode confidence in the whole strategy.

The second problem is that even high-quality proprietary data may not create a durable advantage if foundation model improvements keep democratizing capability across competitors. Interviewee 11 stated it directly: “Most AI-driven roll-up narratives center on cost-side efficiencies. The issue is that these advantages tend to be short-lived. Any company pursuing a durable advantage must assume competitors will eventually deploy similar AI tools, whether off-the-shelf or built in-house.”

There may be real value, however, in the data that has never been digitized, i.e. the information sitting in the physical files and archives of small businesses. Because it has never been digitized, it has never made it into any LLM's training set, which makes it entirely proprietary. As Interviewee 10 observed, “All this offline knowledge is becoming even more valuable now in the age of AI.”

Repeatable Transformation Playbooks Across Multiple Acquisitions

The third assumption is that the AI implementation process can be standardized and repeated across acquisitions with a declining marginal cost. Once the approach is proven on one acquisition, the same tools and processes should in theory deploy across the next ones, each benefiting from the accumulated learnings.

Interviewee 2's three-layer architecture embodies this assumption:

1. Universal Composable Infrastructure (Build Once, Reuse Everywhere): basic functions used repeatedly across all platforms.
2. Vertical-Specific Agents (Build Custom for Each Industry): specific to each industry.
3. Integration and Orchestration (Build Custom Connectors): custom connectors between systems.

The implication is that 60-70% of the technical work is reusable across acquisitions, with 30-40% requiring vertical customization.

Interviewee 3 described the implementation process that supposedly becomes repeatable: “We start with a shadow period, on-site workflow observation. We’re looking at the mix of low-hanging fruit automation opportunities plus where we need custom API builds. The focus is on native end-to-end workflows, around 70% human workflow automation. Our goal is to rip and replace entire workflows, not just add point solutions on top of existing processes.”

Others challenged the repeatability assumption. Many businesses that look identical from the outside run on unique workflows, legacy systems, and organizational cultures. Interviewee 4, who is pursuing this strategy in Germany, noted that even within a single vertical, heterogeneous operations may require a “target-by-target approach” where the product team builds customized modules for each acquisition instead of deploying something standardized.

The change management burden also tends to be heavier than anticipated. Even with ownership control, transforming the way existing employees work and convincing them to adopt AI-augmented processes takes sustained effort. Interviewee 2 again: “Traditional ‘turn the Titanic’ transformations fail without CEO conviction.” Repeatability, in other words, depends as much on finding CEOs who are predisposed to change as it does on the technical architecture, and that is a significant constraint on the universe of acquirable companies.

AI-Generated Efficiencies Create Durable Competitive Moats

The fourth assumption is that margin improvements from AI create a sustainable competitive advantage that compounds over time. While competitors stay trapped in labor-intensive models, the AI-enabled operator can in theory underprice them while maintaining superior margins, gradually taking share and reinforcing the advantage through scale.

Interviewee 2 articulated this view: “Margin gap between AI-enabled and traditional operators widens over time. Traditional PE rollups can pursue SG&A cuts but cannot match agent-driven throughput.”

This is, in my opinion, the assumption that faces the most severe challenge, because cost advantages are transient in an era where AI tools democratize this quickly. One investor was frank about it: “Any company pursuing a durable advantage must assume competitors will eventually deploy similar AI tools. The technology often isn’t proprietary; it’s built on foundation models from providers like OpenAI, Anthropic, and Google that are broadly available. The workflows aren’t secret either; they’re shared across industries and rapidly copied.”

The competitive dynamics follow a predictable pattern. An innovator deploys AI and cuts costs by 50%, capturing the full margin upside at first. Rivals follow and price competition intensifies. Service providers start undercutting each other, prices drop by 30%, then 40%, then further. What started as margin expansion turns into a fight to preserve revenue. Average contract values fall, and the business needs ever more customer acquisition just to keep the top line flat.

Ownership Accelerates Technology Adoption and Overcomes Implementation Resistance

The fifth core assumption is that ownership removes the primary barrier to AI transformation in traditional industries, which is resistance to change. Instead of trying to sell software to skeptical SMB owners, you roll up the businesses and mandate the transformation, which in theory accelerates time to value.

Interviewee 2 emphasizes selection over transformation: “We’re only buying companies where the CEO has already bought into the strategy and actually cares about the strategy. Expected equity rollover: 30–40%. CEO participates directly in change management.” Interviewee 4 added: “We’ve learned that rollover below 20% doesn’t create real alignment. They optimize for short-term earn-outs, not long-term transformation. The sweet spot is 30-40%, enough skin in the game that they care deeply about success.”

This selection-based approach is an admission that ownership alone is not enough. Cultural fit and founder commitment are prerequisites for a successful implementation and not outcomes of one. As one operator put it, “Rather than believing you can force transformation on any business with the right change management process, these operators recognize that cultural fit and founder commitment are prerequisites.”

The ownership assumption also runs into what I've come to think of as the two-company problem. You are building an industry-leading AI software company, and attracting the talent that requires, while simultaneously building an operationally excellent service business, and those two things demand nearly orthogonal skill sets.

Interviewee 2 described the talent side of it: “Most engineers want to build cool tech. But in a roll-up, you need engineers who understand they’re building boring infrastructure that just needs to work reliably at scale. That’s a very different personality type.”

Structural Challenges Undermining the Model

The theoretical logic of AI-enabled roll-ups is appealing. The experience of operators, however, reveals systematic challenges that undermine value creation along several dimensions at once, and understanding these is essential to evaluating when, and whether, the strategy can succeed.

Data Quality Issues: The Hidden Six-Month Tax

The most consistently underestimated challenge is the state of data infrastructure in the companies being acquired. Interviewee 7 again: “Everyone talks about proprietary data as a moat. But honestly? Most SMBs have terrible data. It’s in spreadsheets, inconsistent formats, missing fields. You spend the first six months just getting clean data before you can even think about training models.”

The data quality problem shows up in a number of ways:

Structural Chaos: Customer data scattered across multiple systems that were never integrated. Critical information living in email threads, text messages, or the memory of a single employee instead of a central database. Different departments keeping separate records with no reconciliation process between them.

Format Inconsistency: The same data point recorded in different formats depending on the year or the location. Address fields that contain full addresses in some records, city only in others, and abbreviations in still others. Sometimes AI handles this easily. Sometimes it does not.

Missing Information: Critical fields left blank in 30-50% of records. Optional fields treated as optional even when they hold information the AI implementation depends on. Historical gaps where the data was never collected in the first place or has since been lost.

Quality Degradation: Data entered hastily during busy periods with no quality control afterward. Duplicate records with slight variations, so that nobody knows which is the canonical version. Outdated information that never gets flagged or updated as the customer's situation changes.

Interviewee 3 implicitly acknowledges this timeline when describing their shadow period, which functions partly as data archaeology, i.e. discovering where the information actually lives and how it can be extracted, cleaned, and structured for AI to consume.

The Two-Company Problem: Building Software and Running Services Simultaneously

The second structural challenge comes from the hybrid nature of the model itself. Operators are not building one world-class organization. They are building an industry-leading AI software company and an operationally excellent service business at the same time, and the skills, cultures, and management approaches required for each are very different.

Top AI engineering talent has traditionally gravitated towards pure technology companies with engineering-first cultures, significant equity upside, and intellectually stimulating problems. They

show limited enthusiasm for joining a roll-up of accounting firms, call centers, or staffing agencies, no matter how the AI story is framed (though this may be changing given the elite positioning of some of the leading AI-buyout PE firms). The upside looks bounded compared to the seemingly unlimited potential of a pure technology venture.

Interviewee 2 drew a crucial distinction: “Most engineers want to build cool tech. But in a roll-up, you need engineers who understand they’re building boring infrastructure that just needs to work reliably at scale. That’s a very different personality type.” Successful roll-ups need more than talented engineers. They need engineers who are comfortable with infrastructure work, reliability engineering, and incremental improvement, and who are not looking to do cutting-edge research.

The two-company problem extends beyond recruiting and into organizational design. Traditional PE firms are structured with investment professionals making capital allocation decisions and operating partners executing the post-acquisition improvements. That hierarchy suits AI-enabled roll-ups poorly, because technical judgments about automation feasibility should inform acquisition decisions, and not come after them.

Interviewee 2 addressed this explicitly: “If you’re not tech-native in your approach, it’s a turnoff to the best engineers out there. You need to elevate them within the org. Engineers participate in deal screening, have veto power over acquisitions where the technical transformation thesis seems weak, and receive equity packages competitive with what they’d earn at well-funded startups. The CTO is a true C-level executive, not a director reporting to the COO.”

Creating an engineering-first culture is hard, though, when the bulk of the organization runs service businesses with fundamentally different requirements. A service business succeeds on execution reliability, customer relationship management, and operational consistency. Software development requires experimentation, tolerance for failure, and rapid iteration. Reconciling those two orientations inside a single organization creates a tension that never fully goes away.

Public Market Evidence: AI-Transformed BPO Firms and Persistent Low Multiples

One of the more discouraging data points comes from the public markets. Today's “AI-transformed” BPO firms, companies like Concentrix, Genpact, and Infosys that have invested heavily in automation and AI, trade at 4-15x EV/EBITDA. Their pure software counterparts, Salesforce, ServiceNow, and Workday, command 19-72x.

	Company	EBITDA Margin	EV/EBITDA Multiple	Description
Software-Enabled Services	Concentrix	14%	4.03	Concentrix is a global customer experience (CX) and outsourcing company that designs, builds, and runs customer engagement solutions, including contact centers, back-office processing, and digital transformation services for major brands. **Software-enabled services –** their core business is delivering human-led services powered by proprietary and third-party technology platforms, not selling software subscriptions.
	Genpact	20%	9.49	Genpact is a professional services firm specializing in business process operations, digital transformation, analytics, and AI-driven solutions, originally spun out of GE. **Software-enabled services –** they primarily sell consulting and managed services where software/AI tools enable the outcomes, rather than licensing standalone SaaS products.
	Infosys	24%	14.86	Infosys is one of India's largest IT services and consulting companies, offering application development, maintenance, systems integration, cloud migration, and consulting services worldwide. **Software-enabled services –** classic global IT/BPO services company; they build and run solutions for clients but do not primarily sell subscription software products.
Software	Salesforce	34%	19.73	Salesforce is the world's leading customer relationship management (CRM) platform, providing cloud-based applications for sales, service, marketing, commerce, and analytics. **Pure SaaS company –** their entire business model is subscription-based software subscriptions and platform usage.
	Workday	31%	54.5	Workday offers enterprise cloud applications focused on human capital management (HCM), financial management, payroll, and analytics. **Pure SaaS company –** built from the ground up as a subscription-based software product company with minimal reliance on professional services revenue.
	ServiceNow	21%	71.6	ServiceNow provides a cloud-native platform that automates workflows, IT service management (ITSM), HR, customer service, security operations, and other enterprise functions. **Pure SaaS company –** subscription-based platform with virtually no significant services revenue compared to product subscriptions.

That valuation gap reflects fundamental perceptions about business quality, scalability, and defensibility. For a number of reasons, public investors view services businesses, even AI-enhanced ones, as a different kind of asset from software.

Labor Intensity Concerns: Despite automation, a services business retains a meaningful human labor component, and the market prices in ongoing labor cost inflation and execution risk. A software business, once built, scales at minimal marginal cost. The unit economics diverge dramatically.

Revenue Quality Perceptions: Services revenue requires continuous delivery and performance. Customers churn when service quality slips. Software revenue, particularly SaaS with multi-year contracts, is more predictable and more durable, and the market pays for predictability.

Competitive Moat Skepticism: Services businesses face continuous competition from traditional providers and from new AI-enabled entrants, and public investors question whether an automation advantage can be defended. Software companies with network effects, switching costs, or data moats command premium valuations because their competitive position looks more sustainable.

Growth Profile Differences: Software companies post 30-50%+ annual growth rates that services businesses rarely reach. Even when AI enables capacity expansion, services growth tends to be constrained by implementation bandwidth and geographic expansion, and markets reward growth rates disproportionately.

The implications for exit valuations are sobering. If public markets persistently value AI-transformed services businesses at 8-12x EBITDA, private market exits will likely land at 6-10x. That is a long way from the 10-30x exit multiples that many AI roll-up pitch decks project, and it dramatically constrains the return profile.

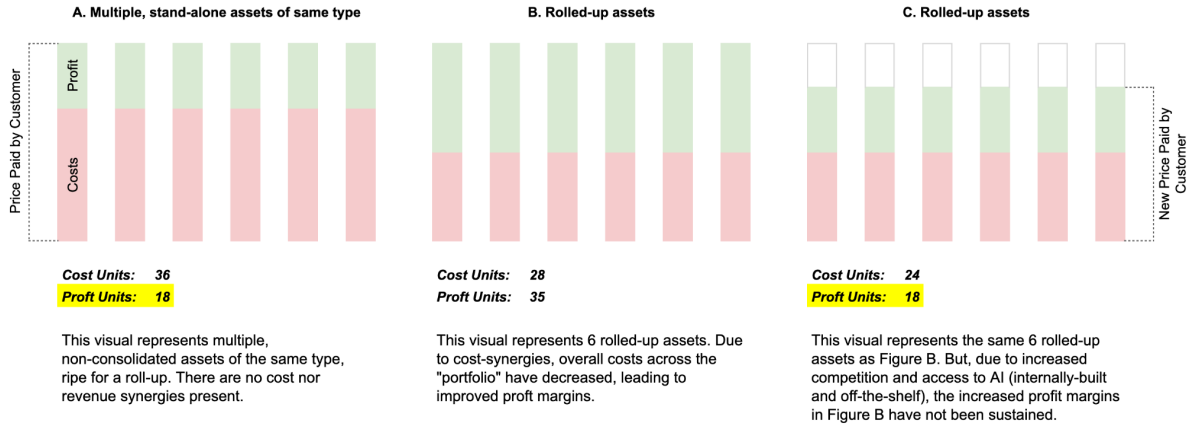
Systematic Margin Erosion as AI Capabilities Democratize

Another challenge is the likely erosion of AI-driven cost advantages over time.

An innovator deploys AI automation and reduces costs by 40-50%. With competitors still running traditional models, the innovator captures the full margin upside. EBITDA margins expand dramatically and the strategy looks validated.

Then the competitors notice. Some build similar capabilities internally. Others buy off-the-shelf tools from vendors who have productized the same workflows. A few get acquired by PE firms pursuing a catch-up strategy of their own.

Once multiple providers operate at similar cost structures, competition shifts to price. The innovator now has to choose between holding price and losing customers or cutting price and sacrificing margin. Most operators cut price to defend share, a price war follows, and prices fall until margins compress back toward the competitive baseline. The value of the innovation ends up accruing to customers through lower prices, and only briefly to the operators through margin.



This dynamic is particularly severe in highly competitive services markets with low switching costs. The pattern is already visible in segments of business process outsourcing, where AI-enabled providers compete their advantages away within 18-24 months of deployment. For a venture-backed operator that needs a five-to-seven year hold to generate an acceptable return, an 18-month advantage window is not enough.

Slow and Messy Implementation Cycles vs. Investor Time Horizons

As discussed, there is a mismatch between AI implementation timelines and investor return expectations. Pitch decks project margin improvement within three to six months of acquisition. The reality involves much longer cycles.

Having spoken with several operators, the following timeline seems closer to the truth:

3-6 months for shadow period and workflow mapping: Observing existing processes, documenting steps, identifying automation opportunities, and building consensus around the changes.

6-9 months for data clean-up and implementation: Resolving data quality issues, building APIs, integrating systems, and developing and testing the automation workflows.

3-6 months for refinement and scaling: Iterating on the initial implementations, addressing edge cases, training employees, and measuring results.

In short, 12-18 months can pass before meaningful margin improvement shows up. In acquisitions with particularly poor data hygiene, 24-plus month implementation cycles seem common.

These timelines create severe pressure inside a venture structure. Investors evaluating a ten-year fund life expect meaningful progress within 24-36 months. When each of several acquisitions needs an 18-month transformation cycle, the math gets difficult.

When cultural resistance, technical obstacles, or organizational skepticism show up, timelines stretch further. Projects planned for 12 months commonly take 18-24, which compounds the pressure.

Absence of Durable Revenue-Side Synergies

Another difficulty is the absence of sustainable revenue synergies. One investor emphasized this: “Revenue-side synergies are far more sustainable. When aggregation via AI can enable companies to produce incremental revenue lift while simultaneously benefiting from AI-powered cost advantages, then you have a compelling and durable profit margin.” In my conversations, most operators struggled to articulate a clear path to revenue synergies through AI.

Cross-sell opportunities are often limited or non-existent. Owning multiple accounting firms gives you no incremental product to sell to the existing customers. Network effects rarely materialize, because one call center does not become more valuable because you own five others. Brand value compounds slowly, if at all, since local service businesses often run under separate brands where regional identity matters more than national recognition.

Worse, some AI implementations actively cannibalize revenue. When a business bills by the hour and AI reduces the hours required, automation reduces the top line even as it improves the margin. Interviewee 5 identified a related problem: “Controller time savings don’t necessarily translate to fireable positions,” and therefore do not necessarily translate to significant cost savings either.

High Change Management Burden Without Committed Founders

A fourth challenge is underestimating change management, particularly when the founders of the acquired companies lack real commitment to the transformation. Interviewee 4 described the lesson: “We’ve learned that rollover below 20% doesn’t create real alignment. They optimize for short-term earn-outs, not long-term transformation. The sweet spot is 30-40%, enough skin in the game that they care deeply about success.”

The change management burden shows up at multiple levels.

Employee resistance is often systematic and, frankly, rational. Workers fear losing their jobs, resist changes to their workflow, and resent technology that gets imposed on them. Even when an operator carefully messages “AI enablement, not replacement,” employees can see where automation is headed and draw their own conclusions about their long-term employment.

Cultural integration compounds the problem. Service businesses often have deeply embedded cultures built around relationship management, personal service, and tacit knowledge. A technology-first approach feels foreign and threatening to them. As Interviewee 1 noted, successful adoption requires “Emphasis on AI enablement, not job replacement. Migration of mundane tasks to agents; humans handle higher-value tasks.” Getting there requires sustained leadership commitment over a long period.

Undercapitalized Transformation Efforts When Debt Is Involved

A final challenge emerges when operators use excessive leverage to finance their acquisitions. Debt improves equity returns in the good scenario and creates severe constraints during the transformation period.

Interviewee 8 explained the problem: “We’re probably 100% equitized on the first two or three businesses for two years. A, because they’re subscale and you can’t have a meaningful conversation with lenders at that size. And B, because you’re going to have that J-curve effect.”

Debt covenants typically require minimum EBITDA thresholds. When an AI implementation temporarily depresses margins, covenant violations get closer, which creates a perverse incentive to abandon the transformation mid-implementation, which is exactly the moment when persistence would generate the most value. Interviewee 2 gave his rationale for running with zero leverage: “No debt, no forced cost cuts. We avoid dependence on lenders during early operational disruption.”

In practice, the combination of unpredictable implementation timelines, uncertain margin improvement, and rigid debt covenants makes the whole thing fragile. When the transformation takes 18 months instead of 12, or delivers 25% margin improvement instead of 40%, a leveraged operator faces an existential problem. An all-equity approach gives up some return efficiency and in exchange preserves the option to keep going through the obstacles.

Conditions Under Which the Strategy Can Work

Most AI-enabled roll-up ventures face systematic headwinds. Certain conditions make the strategy more feasible, and those conditions look a lot like the conditions that make any roll-up strategy work. Understanding them matters both for operators choosing a vertical and for investors evaluating an opportunity.

Vertical Characteristics: Recurring Revenue, Long-Term Contracts, and Low Churn

The foundational requirement is revenue predictability. Interviewee 8 was explicit: “We’re not doing deals with high churn. When we implement AI automation, it’s not always 100% smooth, so we need the benefit of at least six months to a year to implement these changes.”

Revenue predictability shows up through a few mechanisms.

Multi-year contracts create a time horizon long enough for the transformation benefits to materialize. When customers commit to two or three year engagements, implementation disruption is tolerable. Single-year contracts provide too little runway, and month-to-month agreements are nearly impossible to manage during an active transformation.

Low customer churn enables compounding value creation, because each automation improvement benefits the same customer base over an extended period. High churn forces continuous customer acquisition, and the automation benefits end up accruing to customers the operator never keeps long enough to justify the implementation investment.

Recurring revenue models provide cash flow stability during the transformation. Monthly recurring revenue supports debt financing where appropriate, provides runway for experimentation, and creates some forgiveness for implementation stumbles. Project-based revenue creates existential pressure, where any disruption to implementation threatens the viability of the business itself.

Tasks Well-Suited for Multi-Stage Automation

Not all white-collar labor is equally automatable. Interviewee 1 offered a useful categorization:

1. Customer support and FAQ handling.
2. Document creation (presentations, marketing, emails).
3. Repetitive data tasks such as forms, insurance checklists.
4. Reasoning tasks such as underwriting, legal Q&A.

The automation potential varies dramatically by category.

Category 1 (Customer Support/FAQ) offers 70-90% automation potential with current technology. LLMs excel at answering common questions, routing inquiries, and handling routine interactions.

These tasks are highly repetitive, well documented, and low risk when the AI occasionally gets something wrong. Implementation timelines run from weeks to months.

Category 2 (Document Creation) provides 60-80% automation potential. AI is strong at generating written content that previously consumed significant human time, and this category often yields the quick wins that build organizational confidence. Quality control and brand voice consistency still require human oversight.

Category 3 (Repetitive Data Tasks) delivers 70-85% automation potential. Forms processing, data entry, and basic reconciliation are tedious for humans and straightforward for a properly trained AI system, especially one combined with OCR and a structured database. These implementations create substantial margin improvement.

Category 4 (Reasoning Tasks) currently offers only 30-50% automation potential. Underwriting decisions, legal analysis, and complex judgment require domain expertise and the ability to explain the reasoning. AI augmentation makes more sense here than full automation. These tasks create the highest value and resist complete automation.

Successful implementations typically start with Categories 1 and 2 to build confidence, move to Category 3 for the substantial margin improvement, and then approach Category 4 with hybrid human-AI workflows. Interviewee 1 observes, “Bookkeeping businesses have enormous Category 3 surface area (80%+ automation potential). Call centers are dominated by Categories 1-2 (50-70% automation potential). MSPs have meaningful opportunities across all four categories but with more complexity in each (40% automation potential overall).”

The implication is that vertical selection should explicitly evaluate the mix of automation categories present in the core workflows. Businesses heavy in Categories 1-3 offer superior near-term opportunities. Businesses dominated by Category 4 work need a longer horizon and more sophisticated technical capability.

Market Dynamics: Fragmentation, Consolidation Tailwinds, Rising Cost Structures

Beyond task-level automation potential, broader market dynamics heavily influence whether the strategy is viable. Several characteristics signal favorable conditions, and none of them is unique to the AI version of the approach. They are table stakes for any roll-up.

Industry Fragmentation: Thousands of potentially acquirable assets within the target size range.

Active Consolidation Already Underway: Markets already experiencing PE-driven consolidation offer natural exit pathways. Strategic and financial buyers understand the vertical, valuations are more transparent, and comparable transactions provide pricing benchmarks. Being early to a consolidation trend can also be beneficial, however.

Aging Ownership Demographics: Sectors where baby boomer owners are approaching retirement offer a structural advantage. Succession-focused sellers often prioritize deal certainty

and a smooth transition over the maximum price, and they tend to respond well to a buyer offering a partnership instead of pure financial extraction.

Rising Cost Structures: Industries facing systematic cost inflation create a tailwind for AI-driven efficiency. When labor costs rise 5-10% annually, AI's relative cost advantage compounds every year.

These market dynamics are largely outside the operator's control and they dramatically influence outcomes. Entering a favorable market at the right time can offset execution problems. Exceptional execution in an unfavorable market rarely generates an acceptable return.

Path to Revenue Expansion Rather Than Pure Cost Reduction

The single most important condition for success, in my view, is a clear path to revenue expansion beyond pure cost reduction. As one investor emphasized: “Revenue-side synergies are far more sustainable.”

Revenue expansion shows up through several mechanisms:

Capacity Expansion Without Proportional Cost Increase: When AI enables the same team to serve more customers, revenue can grow 50-100% while costs grow only 20-30%. Interviewee 3 described this operating leverage explicitly: “AI enables service capacity expansion vs headcount reduction. More and elevated client-facing time can equate to higher pricing.” That model seems far more sustainable to me than margin expansion through cost cuts alone.

Quality and Speed Improvements Enabling Premium Pricing: When AI materially improves service quality or delivery speed, customers pay a premium. Traditional service businesses usually face a tradeoff between quality, speed, and price. AI enables simultaneous improvement across all of them, which justifies premium positioning.

Product Expansion Into Higher-Value Services: AI sometimes makes it possible to offer entirely new services that were previously uneconomic. Accounting firms, for example, have historically avoided proactive advisory work because of how labor-intensive it is. AI-enabled financial analysis and scenario modeling makes advisory viable, which expands both the addressable market and the average revenue per customer.

Scale-Dependent Strategic Advantages: At sufficient scale, an aggregated platform captures advantages that go beyond the efficiency of any individual business. Interviewee 1 gave the clearest articulation: “Roll-up as wedge into larger prize. At sufficient scale, an aggregated insurance agency platform might monopolize retail insurance distribution, capturing value across the entire value chain.” That kind of strategic option value turns the investment from financial engineering into something closer to real business model innovation.

The contrast between a cost-reduction strategy and a revenue-expansion strategy is decisive. Cost-reduction strategies face systematic margin erosion as competitors adopt the same tools. Revenue-expansion strategies create defensible advantages through quality differentiation,

capacity expansion, and market position. Operators have to be honest with themselves about whether their vertical and their approach offer real revenue upside or rely mostly on cost arbitrage.

Ability to Capture Increasing Returns to Scale Across the Portfolio

A final condition for success is the ability to demonstrate increasing returns to scale, where each incremental acquisition improves unit economics for the entire portfolio and not just for itself.

True economies of scale appear when:

Shared infrastructure costs decline per business as the portfolio grows. A $1M investment in composable AI infrastructure is far more valuable serving twenty businesses than one. Data network effects materialize where a larger dataset improves model performance across every portfolio company. This requires real workflow similarity, though. Acquiring dissimilar businesses produces data volume without any quality improvement.

Negotiating leverage increases with scale. Larger platforms command better terms from suppliers, technology vendors, and distribution partners. Talent leverage improves because a specialized engineer can support multiple businesses, which again requires a standardized technical architecture. Highly customized implementations eliminate the talent leverage entirely.

Interviewee 2's composable infrastructure is the embodiment of this principle. Build the reusable technical layers once, then deploy them across the portfolio at minimal marginal cost. It only works when the acquired businesses share enough workflow similarity to benefit from the standardized tooling.

Playbook: Frameworks and Best Practices from Operators

The operators pursuing AI-enabled roll-ups with discipline have developed frameworks and practices worth examining. They are not universal rules, but they do separate the approaches that are working from the ones that are struggling.

Rule One: Revenue Synergies Over Cost Efficiencies

As multiple investors put it, “When aggregation via AI can enable companies to produce incremental revenue lift while simultaneously benefiting from AI-powered cost advantages, then you have a compelling and durable profit margin.”

Ideally, revenue synergies demonstrate increasing returns to scale, where each incremental acquisition improves the unit economics of the whole portfolio. That might show up as cross-sell from a broader product suite, higher commission rates or better supplier terms that are only available above a volume threshold, network effects where each additional node makes the platform more valuable for everyone on it, or brand value that compounds across locations.

In my interviews, investors and operators often admitted they had no clear path to revenue synergies through AI. That admission signals a fundamental weakness in the strategy.

Rule Two: Roll-Up as Wedge Into Larger Prize

The most successful AI roll-ups treat consolidation as the entry point to a more valuable strategic position and not as the end game.

At sufficient scale, an aggregated platform occupies a position from which it can capture value across the entire value chain. An insurance agency roll-up that consolidates retail distribution might eventually negotiate directly with carriers, disintermediating the wholesalers and capturing margin across the stack. A healthcare staffing platform at scale might integrate backward into training and certification, forward into workforce management software, and laterally into adjacent healthcare services.

This strategic optionality is what turns the investment from pure financial engineering into business model innovation. The roll-up phase functions as patient capital deployment that builds market position, and the ultimate prize is strategic leverage that smaller players simply cannot access.

Rule Three: Speed-to-Close as Competitive Advantage

Interviewee 4 articulated this clearly: “In traditional M&A, you’re looking at 6-9 months to close. We can close in 60-90 days because we’re not doing debt diligence, we don’t need bank approvals, and founders appreciate the certainty.”

The advantage expresses itself in several ways. An all-equity buyer eliminates deals falling apart during debt diligence. Transformation can begin immediately after close. Sellers who have been burned by drawn-out negotiations remember who closed quickly, which turns into proprietary deal flow. There are no covenant constraints on operational flexibility, and closing 90 days faster lets you offer a slightly lower price while remaining the more attractive buyer.

This advantage requires an equity-heavy capital structure, though. Operators have to choose between capital efficiency through leverage and speed and certainty through equity. Optimizing for both at once is not possible.

Capitalization Structure

The consensus view is that traditional venture structures are the wrong fit. Finding the right alternative remains hard.

Advantages of Equity-Heavy Early Phases

There are several advantages to over-equitizing the initial acquisitions. Interviewee 2 stressed strategic flexibility: “Our platforms are not putting any debt on. All equity. We candidly will invest a little bit in R&D in each of these companies… So we will dip into burn to start off with.” The willingness to accept short-term margin compression in exchange for long-term improvement is

crucial. Debt covenants with minimum EBITDA requirements force you to abandon AI initiatives mid-implementation, at exactly the point where persistence pays off most.

Interviewee 8 synthesized both perspectives: “We’re probably 100% equitized on the first two or three businesses for two years. A, because they’re subscale and you can’t have meaningful conversation with lenders at that size. And B, because you’re going to have that J-curve effect.” Equity-heavy structures particularly suit early-stage platforms, before scale and predictability have arrived.

Drawbacks of VC-Style Structures

Traditional venture structures impose maximum dilution at the point of minimum value creation, distribute gains proportionally from dollar one instead of after a hurdle, and give founders an incentive to maximize the capital raised regardless of what the dilution costs them.

Traditional VC fund lives are also misaligned with roll-up timelines. Interviewee 2 again: “The traditional fund structure doesn’t really work for this, especially when you’re trying to do roll-ups for multiple different industries.” A ten-year fund life with return pressure in years seven and eight clashes with a strategy that needs 12-15 years.

Role of Leverage Once Scale Is Achieved

The equity-versus-debt debate does not have to be binary. The optimal approach is often an equity-heavy early phase that transitions to leverage once scale, predictability, and a proven transformation capability are in place.

Interviewee 8 described this sequencing: “Bridge to debt at ~$10M EBITDA milestone.” At that scale, lenders get comfortable underwriting the business, the transformation playbook has been validated, and covenant constraints become manageable. Introducing leverage at that point improves equity returns without giving up strategic flexibility.

The transition requires discipline, though. Many operators, having reached scale with equity, reflexively pursue maximum leverage to “monetize” the success. That often backfires when the next acquisitions turn out messier than expected and covenant pressure forces suboptimal decisions.

HoldCo Structure as Alignment Mechanism

Interviewee 2's advocacy for a HoldCo structure addresses many of the fund structure misalignments: “No fund timeline pressure. Flexible capital deployment across industries. Simplified governance. Potential tax efficiency.” A HoldCo requires a different kind of investor, however, such as family offices, ultra-high-net-worth individuals, and other flexible allocators.

The HoldCo structure particularly suits operators who plan to hold assets long term, potentially permanently. Instead of optimizing for a five-to-seven year exit, a HoldCo enables indefinite ownership with distributions from cash flow. That aligns beautifully with AI roll-up economics, where value creation compounds over decades and does not need to crystallize in a single exit event.

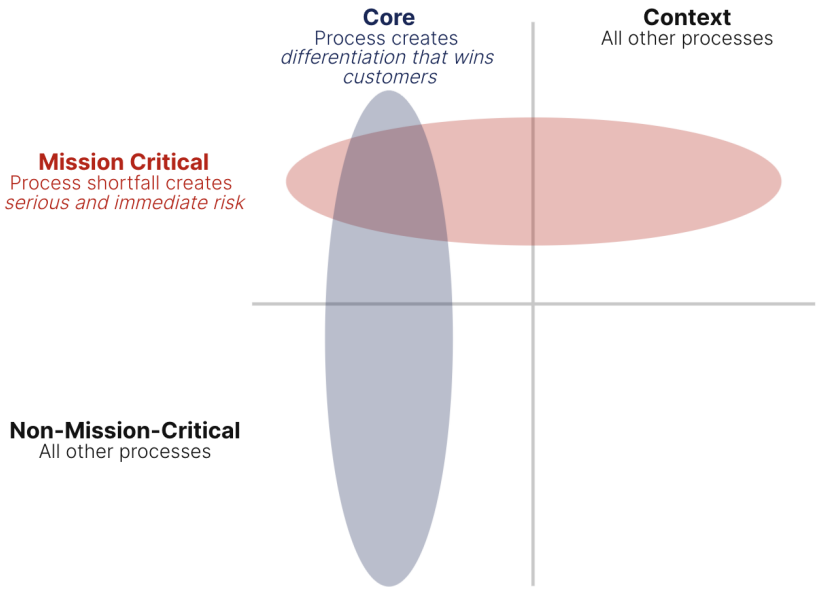
A HoldCo also requires operators who are comfortable with concentrated exposure and investors who are willing to accept illiquidity. The structure trades diversification and liquidity for alignment and control, which appeals to certain operator and investor profiles and not to others.

Build-vs-Buy: Core vs Context Framework

One of the most important decisions is which software to build internally and which to buy off the shelf. I find it helpful to break tasks and processes into core functions and context functions.

Core functions directly drive competitive advantage and define why the company wins. Context functions are necessary, and they do not create differentiation.

Applied to AI roll-ups: build custom AI-driven tools for the core functions that strengthen your competitive advantage, and buy standardized solutions for the context areas where commoditization has already happened.



Build Only Where It Drives Differentiation

Build custom AI when the function directly determines why customers choose you over competitors and when excellence in it creates measurable revenue lift or margin improvement. Build it when you have proprietary data that enables differentiated model performance. If the workflow is complex, specific to your business, requires deep integration, and you'll deploy it across enough businesses, at least fifteen or so, to justify the development cost, then it makes sense. The capability compounds with scale.

The point is that building custom software is expensive, $500K to $1M-plus for a sophisticated automation platform. That investment only makes sense when it gets deployed across multiple

acquisitions and when the capability really differentiates. For most functions, the off-the-shelf solution is the better choice.

Buy for Commoditized Layers

Buy off the shelf when the function is necessary and does not influence competitive position, when excellent vendors with mature products already exist, when the capability is standardized across industries, when speed to implementation matters more than a perfect fit, and when building would distract from the core development work.

In practice, 70-80% of software needs should be met through vendors: CRM, accounting, HRIS, project management, communication tools. Only the 20-30% that truly differentiates warrants custom development. Correctly identifying which 20-30% that is requires deep technical sophistication, and that is where a lot of teams go wrong.

Engineering Talent Model

Attracting and retaining top engineering talent is essential and systematically difficult for an AI roll-up. Interviewee 2 described the fundamental requirement: “If you’re not tech-native in your approach, it’s a turnoff to the best engineers out there. You need to elevate them within the org.”

Treat Engineers as Strategic Equals

The traditional PE hierarchy positions engineers as support staff reporting into operations. That structure is fatal for an AI-enabled roll-up, where technical judgment should inform strategic decisions.

Interviewee 2 described the required shift: “Engineers participate in deal screening, have veto power over acquisitions where the technical transformation thesis seems weak, and receive equity packages competitive with what they’d earn at well-funded startups. The CTO is a true C-level executive, not a director reporting to the COO.”

That elevation is costly in traditional PE terms. Engineering compensation typically runs 50-70% of investment professional compensation inside a PE firm, so parity means roughly doubling the engineering budget. Without that investment, though, attracting talent capable of executing a sophisticated AI implementation is incredibly difficult.

Perhaps the most radical departure from traditional PE is giving engineers formal veto authority over acquisitions. The rationale is sound. Deal professionals often overestimate implementation feasibility, underestimate data quality problems, and fail to recognize technical complexity. An engineering veto forces an honest assessment before capital gets committed, and it is better to pass on a marginal deal than to close an acquisition that turns out to be untransformable.

That authority requires engineers with business judgment, though, and not just technical capability. Engineers who understand both technology feasibility and business economics are exceptionally hard to find. Many operators compromise by requiring engineering approval instead of granting a pure veto, which creates a consultation obligation without absolute authority.

Where the Strategy Makes the Most Sense

Not all service industries offer the same opportunity for an AI-enabled roll-up. My interviews with operators surfaced specific characteristics that separate the promising verticals from the unpromising ones.

Verticals with Rich Category 2-3 Workflow Exposure

Interviewee 1's automation category framework, discussed earlier, is the most useful guide here.

Verticals dominated by Categories 2 and 3 offer superior near-term potential. Document-heavy workflows in legal services, accounting, consulting, and insurance, where a large share of time goes into creating presentations, reports, contracts, and correspondence, are highly automatable. Data-intensive processes in financial services, healthcare administration, and compliance, which involve forms processing, data entry, reconciliation, and reporting, offer substantial margin improvement.

Verticals dominated by Category 4, i.e. complex judgment and reasoning, need a longer horizon and more sophisticated technical capability. AI augmentation adds value there, but reaching the 30-40% margin expansion that justifies the strategy is considerably harder.

Industries with Natural Structural Revenue Growth

Several operators emphasized looking for verticals where revenue grows on its own, independent of customer acquisition. That characteristic creates a tailwind that compounds with AI-driven margin expansion.

Examples include insurance services, where premiums rise with inflation and revenue from existing customers grows without any additional sales effort. Healthcare-related services, where aging demographics and medical cost inflation drive systematic revenue growth. Regulatory compliance services, where complexity increases over time and expands both the scope and the pricing of existing engagements. Professional services to growing companies, where the customer's success automatically expands the engagement.

Structural revenue growth is particularly valuable because it lets the operator focus on retention and margin improvement instead of constant customer acquisition. When a customer who generates $10,000 in year one organically grows to $13,000 by year three purely from inflation and scope expansion, the compounding with AI margin improvement is powerful.

Businesses with Workflow Homogeneity Supporting Compounding Automation

Interviewee 2's emphasis on reusable infrastructure depends fundamentally on workflow similarity across acquisitions. When the acquired companies perform essentially identical tasks with minor variations, the automation investment compounds. When the workflows differ substantially, every acquisition requires custom development.

Positive indicators include:

Industry standardization: Regulatory requirements that force similar processes across competitors, as in financial services compliance or healthcare documentation.

Common software platforms: Verticals where most companies run on the same underlying systems, which enables a standardized integration approach. Accounting firms on QuickBooks, staffing agencies on a handful of common applicant tracking systems.

Workflow convergence: Industries where best practices have emerged and been widely adopted, producing procedural similarity.

Negative indicators include:

High customization requirements: Services where each customer demands a unique workflow, such as custom software development or bespoke consulting.

Fragmented tooling: Verticals where every company uses different software with no standardization.

Craft-based work: Services where the individual operator's style defines the offering, such as graphic design or executive coaching.

Sectors Where Incumbents Lag in Tech Adoption and Data Hygiene Can Be Fixed

The ideal verticals combine two characteristics: a meaningful tech adoption lag, where incumbents trail the available technology by three to five-plus years, and fixable data problems, where poor data quality comes from neglect and not from some fundamental impossibility.

Several sectors fit this profile:

Accounting and tax preparation: Incumbents often run on paper-heavy processes and fragmented digital tools despite sophisticated software being readily available. The underlying data, however, exists in structured formats such as general ledgers and tax forms, which makes consolidation and cleaning feasible.

Staffing and recruiting: Many agencies track candidates in spreadsheets or basic databases despite the availability of sophisticated applicant tracking systems. The candidate information exists and just needs consolidation and structure.

Insurance agencies: Independent agencies frequently run on minimal digital infrastructure despite industry-specific platforms offering comprehensive automation. Policy information already lives in the carriers' systems, so the job is extraction and integration and not creation.

Conclusion

The AI-enabled roll-up is neither the pure technology play the optimists describe nor the impossible task the skeptics dismiss. It sits in a complicated middle ground that requires exceptional execution across several dimensions at the same time.

A Third Category Requiring Dual Mastery

The most important insight from this work is that a successful AI roll-up cannot be understood as a technology investment or as a services investment. It is a distinct third category that requires world-class capability in both domains. The operators who succeed will not be technologists learning services operations, and they will not be services operators adopting technology. They will be the rare teams that bring real expertise to both from day one. That requirement narrows the pool of viable players dramatically, and it also suggests that success, where it happens, may be more defensible than either a pure technology or a pure services strategy.

The Execution Gap Is the Strategy

The persistent theme across my operator interviews was not uncertainty about whether AI can transform services businesses. That question is settled. The theme was the size of the gap between the conceptual model and the implementation reality. Data requires six to twelve months of cleaning before meaningful AI work begins. Cultural transformation takes 18-24 months despite ownership control. Integration complexity grows non-linearly with each acquisition.

These are not temporary barriers that better tools will eliminate. They are inherent features of the strategy. Paradoxically, that execution difficulty may be the most durable competitive advantage available. As AI capabilities democratize and cost advantages erode, superior execution across acquisitions, integrations, and transformations becomes the primary differentiator. Being best in class at adopting best-in-class software, especially for context functions, will likely be an edge as well.

A Narrowing Window and Strategic Implications

The opportunity window for AI-enabled roll-ups appears to be time-bound in a way that deserves explicit acknowledgment. As AI tools democratize, the margin expansion thesis weakens, because incumbents will adopt the same automation without any change in ownership. The highest-value opportunities therefore exist in markets where the gap between current operations and what AI makes possible remains widest, and where that gap can be closed faster than competitors can adopt the same capabilities. The implication, I believe, is that speed of execution matters far more than perfect vertical selection. A team that can transform three companies in 24 months in a good vertical will likely outperform a team that spends 24 months picking the perfect one.

Success Requires Ruthless Selectivity

For the subset of operators who meet all of the conditions, i.e. the right vertical, an aligned capital structure, real technical and operational sophistication, and realistic timelines, an AI-enabled roll-up can generate substantial value. The strategy will not work for most teams, most verticals, or most capital structures, and that is precisely what makes it interesting. The opportunity remains relatively uncrowded compared to traditional venture or PE strategies because the barrier to entry is not capital. It is the rare combination of skills and structural alignment that execution requires. Those who can honestly assess whether they have these prerequisites, instead of convincing themselves they can pick them up along the way, have found the true first step in the strategy.