L&T Technology Services

The PE Value Creation Mandate

How the agentic AI cycle is changing the PE software thesis and what value creation teams should do about it

Purposeful.
Agile.
Innovation.

Table of Contents

THE VALUE CREATION INFLECTION: WHY THIS CYCLE IS DIFFERENT 3

The new entry, hold, and exit lens 4

THE EVOLVED VALUE CREATION PLAYBOOK 4

1. Two-speed operating model: run both engines on purpose 4

2. Pricing architecture: from seats to outcomes 5

3. The AI factory: an operating capability, not a project 5

4. KPI architecture: three layers, reviewed monthly 6

5. Commercial engine redesign: pay for what you want sold 6

WHAT SPONSORS CAN ONLY DO: PORTFOLIO LEVEL LEVERAGE 7

Platform-level AI advisory forum 7

Shared GCC strategy: from labor arbitrage to AI leverage 7

Programmatic bolt-ons for agentic capability 7

Cross-portfolio CEO bench and talent benchmarking 8

Shared infrastructure, shared learnings..8

THE OPERATING PARTNER MANDATE 8

What the operating partner owns 8

A 12-month operating calendar for the VCP 9

SUMMARY: THE PLAYBOOK IS NOT NOVEL – THE EXECUTION IS 9

What changes in the underwriting model 10

THE VALUE CREATION INFLECTION: WHY THIS CYCLE IS DIFFERENT

Every platform shift in enterprise software has followed the same pattern. The client/server system displaced the mainframe, the internet took the place of packaged software, and the cloud and mobile displaced on premises. In each cycle, multiples compressed for incumbents who hesitated, category leaders emerged within 24–60 months, and the survivors executed the same moves: stay closer to the customer than the insurgent can, use installed-base scale as a moat the insurgent cannot replicate at speed, and reinvest the savings into the new platform.

What is different in the agentic cycle is not the pattern, but the speed at which the change is underway. While cloud took roughly a decade to reach mainstream enterprise adoption, Agentic AI products are landing in production within quarters. Take the example of Anthropic which reported an annualized revenue run-rate of over $30 billion as of April 2026[1], an exponential growth from $9 billion in revenue at the end of 2025. The global technology industry has never seen this kind of growth with any product in the history

The global software sector has endured nine downturns of meaningful magnitude since 2014, with the current EV/Sales compression sitting well within the historical range. It most closely resembles the 2014 and 2016 dislocations – both of which gave way to extended expansion. Again, the PortCos that emerged stronger from the 2001 and 2008 scenarios were not the ones with the cleanest tech stack, but rather, the ones who stayed closest to their existing customers, took price for the value they delivered, and used the disruption as cover to take cost out.

What is different this cycle is the simultaneity: the speed of the cost-curve collapse and the erosion of switching costs are happening in the same window**. The defensive moat, “we are deeply embedded”, has narrowed. The offensive moat, “we own the customer relationship and the data”, has widened.** The Value Creation teams of PE firms and their portfolio companies will have to exploit that asymmetry before a new entrant does.

Old World Value Creation	What is Breaking in the AI Era?
Switching costs in embedded workflow categories (ERP, CRM, HCM, ITSM, vertical SaaS) are near-permanent.	A small team can build a plausible agentic replacement in ~18 months; the cost curve has collapsed.
New-logo velocity is the primary growth lever to get to exit.	Buyer scrutiny has intensified; NRR is now the most leveraged metric in the SaaS P&L.

[1] https://venturebeat.com/technology/anthropic-says-it-hit-a-30-billion-revenue-run-rate-after-crazy-80x-growth

Old World Value Creation	What is Breaking in the AI Era?
Operating margin expands linearly across the hold.	Funding AI initiatives compresses margin first, then expands. The trajectory is non-linear.
The next buyer prices ARR.	The next buyer prices ARR mix, i.e. legacy vs. AI-led vs. agentic-outcomes revenue at different multiples.

The new entry, hold, and exit lens

At entry. AI-native readiness is now a first-line diligence question, not a technology side-session. The value creation team should underwrite the credible path from current product to AI-native architecture (not a copilot overlay); the quality and uniqueness of the proprietary data a new entrant could not replicate; the optionality in the pricing architecture to shift to outcomes-based model without a disruptive repricing event; and the CEO's belief system. A CEO who is defensive about AI is a hold-period problem.

During the hold. The value creation program reorganizes around a two-speed operating model. Margin becomes a configurable variable. The KPI dashboard expands from the traditional SaaS view to a three-layer view. Board cadence moves monthly, not quarterly.

At exit. The next buyer, whether strategic, sponsor, or public-market, will underwrite AI-native mix as a leading indicator. The exit narrative must separate legacy ARR from AI-led ARR and show a credible line of sight from the former converting into the latter. Sponsors who run this dual-track narrative from mid-hold onward will price better than those that reveal it in the sale process.

THE EVOLVED VALUE CREATION PLAYBOOK

The playbook has FIVE programmatic moves. Each is a Value Creation pillar jointly owned by the PE operating partner and the PortCo CEO.

1. Two-speed operating model: run both engines on purpose

Engine 1: Defend and monetize the installed base. In a cycle where new-logo acquisition costs are rising and buyer scrutiny has intensified; the math of growth has flipped. Every hundred basis points of NRR is worth multiples more in enterprise value than the equivalent new-logo growth because it requires no incremental sales-and-marketing spend and it compounds. The mandate is mechanical: AI-led account expansion plays, contracted-but-undeployed seats activated, ruthless churn forensics. The headroom in the existing book is the cheapest growth available and the growth most insulated from a slower buyer environment.

Engine 2: Build the AI-native product layer. This is not a copilot bolt-on. The bolt-on era is over. Buyers now distinguish between vendors who have wrapped a model around a legacy workflow and vendors whose product is genuinely re-architected around agents that act. The former are being repriced. The latter command pricing power. Engine 2 must be staffed as a separate team, with a separate roadmap, a separate gross-margin profile during build, and direct CEO sponsorship. Retrofitting the AI-native layer through the existing roadmap committee is the most common failure mode observed across the portfolio.

Engine 1 funds Engine 2. Cash and talent flow in one direction. PE-backed software is uniquely suited to this architecture because the sponsor can authorize the cross-subsidy in ways public-market CEOs cannot: the quarterly Street pressure that punishes margin compression for reinvestment is absent. The PortCos that outperform in the 2027–2029 exit window will compress operating margin 100-200 bps in the next four quarters to fund Engine 2, then expand it back 200-400 bps as the AI factory matures. That is the new margin trajectory the VCP team should underwrite and defend to the LP base.

2. Pricing architecture: from seats to outcomes

Traditional SaaS monetizes access. The agentic era monetizes work, e.g., claims processed, tickets resolved, documents generated, code shipped, etc. Vertical SaaS leaders are converging on an outcome-oriented pricing because it gives buyers a familiar headcount-replacement framework and gives vendors a path to revenue that scales with executed work.

The pricing redesign is a value creation project, not a product pricing committee topic. We need to build the architecture now. Do not let the PortCo be forced into a disruptive repricing event after the product is ready and the commercial model is not.

3. The AI factory: an operating capability, not a project

“AI transformation” is misleading because it implies a destination. What the leading PE-backed software companies are building is an AI factory, i.e., a permanent operating capability that continuously applies AI across the product, the go-to-market motion, and the back office.

Three principles anchor this factory:

- **Build AI-native products, not feature overlays**. Proprietary context becomes the moat, continuous learning becomes the roadmap, demonstrable ROI becomes the distribution motion. Pricing shifts from seats and licenses to outcomes and throughput.

- **AI in the SDLC is non-negotiable.** Companies that have moved aggressively are reporting significant reductions in feature-delivery cycle times to the tune of 30-50%, while the ones that have not are losing engineering talent to rivals.

- **AI in operations follows the same logic.** Rip and replace the internal CRM, HCM, project planning, ITSM, and support stacks with agentic alternatives. Publish the productivity metrics. The buyer's first question in 2026 is no longer "does it work?" It is "does it work for you?" The sponsor's standing question at every board should be: which function did we re-architect this quarter, and what was the FTE-equivalent saving? Those savings fund Engine 2 without diluting the equity story.

4. KPI architecture: three layers, reviewed monthly

Running an AI factory requires a KPI framework that goes beyond the traditional SaaS dashboard. Three layers, reviewed at every board:

Layer	What it measures	Representative metrics
Business fundamentals	The investor case	Revenue growth, NRR, GRR, Rule of 40, CAC payback, gross margin
AI-specific	The AI factory's output	% revenue from AI-led products, AI feature attach rate, agentic ARR mix, engineering cycle time, support deflection, sales-cycle compression
Cultural	The AI factory's health	% employees using internal AI tools weekly, training hours per FTE, ecosystem partner count

5. Commercial engine redesign: pay for what you want sold

When Accenture made the digital pivot in the mid-2010s, they capped commission accelerators on legacy services and made digital and cloud sales non-linear. Within four years, "the New" was most of the revenue.

The lesson generalizes. People sell what they are paid to sell. If the sales team is compensated the same on the renewal of a 2019-architecture product as on a 2026-architecture agentic product, the mix will not move. Cap accelerators on the legacy SKUs. Make the agentic SKUs non-linear. The comp plan is the single most underused lever in the value creation toolkit.

WHAT SPONSORS CAN ONLY DO: PORTFOLIO LEVEL LEVERAGE

The PortCo CEO can run most of the value creation playbook – plays below are reserved for the sponsor and the value creation function because they only produce a return at portfolio scale. This is where the operating partner role compounds against public-market incumbents and against sponsors who are running VCPs asset-by-asset.

Platform-level AI advisory forum

Three to five external practitioners or operators who have shipped agentic products at scale, not consultants, meeting quarterly with PortCo CEOs and CTOs. The cost is trivial, but the improvement in signal-to-noise on technical and architectural decisions is substantial. What we have is a scenario where the strongest sponsors are running these at the platform level and sharing the output across the portfolio.

Shared GCC strategy: from labor arbitrage to AI leverage

The labor-arbitrage case for India captives, and increasingly Latin America and Southeast Asia, has been augmented by an AI-leverage case. An engineering team in Bengaluru or São Paulo with full agentic tooling now produces at the cost-adjusted output level that previously required US-coastal teams.

For the portfolio, the question is no longer “should we?” but “which geography, what scope, when?” and the sponsor can stand up shared-services GCC capacity across multiple PortCos at a fraction of the per-asset cost.

Programmatic bolt-ons for agentic capability

The 2021–2023 software M&A pause was a function of valuation gaps. Those gaps are closing in the right direction for buyers of AI-native capability. Three thesis are worth underwriting now across the portfolio:

- **Talent-and-architecture acqui-hires** of 10–40-person teams that have built genuine agentic IP.
- **Tuck-in product acquisitions** that close an Engine 2 capability gap faster than build.
- **Larger combinations** for PortCos on burning platforms consolidating two sub-scale assets into a credible AI-native challenger.

Run this as a programmatic motion with a shared deal team, shared integration playbook, and shared post-close operating cadence. That is an advantage no public-market incumbent can replicate at the speed a sponsor can.

Cross-portfolio CEO bench and talent benchmarking

The CEOs executing this playbook well are leading it personally. Sponsors that have built a CEO bench with an explicit AI-native operator profile, and that are benchmarking AI KPIs across their portfolio, are adding a cycle-defining edge to the operating partner function. The next CEO you appoint should be evaluated on belief as much as on track record.

Shared infrastructure, shared learnings

Agent infrastructure, evaluation harnesses, data pipelines, governance frameworks, model-vendor MSAs, and responsible-AI playbooks are high-fixed-cost functions where the portfolio can amortize across assets. Build the shared-services stack once. Deploy it across the portfolio. Benchmark the results. This is infrastructure the PortCo CEO cannot justify as a single-asset investment but that the sponsor can justify as a platform capability.

THE OPERATING PARTNER MANDATE

What the operating partner owns

- **Discipline.** Monthly operating reviews, not quarterly. A single source of truth for the three-layer KPI dashboard. The discipline that made the best public software companies durable through the earlier resets is the same discipline that will make a PE-backed company defensible through the AI transition.

- **Organizational shape.** Every layer between the buyer and the engineer slows the feedback loop on which AI-native iteration depends. The fastest PortCos have removed two-to-three management layers in the last eighteen months and pushed product and pricing authority down to the segment-leader level.

- **Incentive redesign.** Cap accelerators on legacy SKUs; make agentic SKUs non-linear. The single most underused lever in the VCP toolkit.

- **Uncomfortable questions, repeatedly.** Why are we not using AI in this function? Are our people trained and do we have continuous learning paths, not one-off workshops? Do I have leaders on the CEO's team who actually believe the AI thesis, or am I being managed? The operating partner who cannot answer these affirmatively at the next board has just identified the agenda for the following quarter.

- **Visibility.** The CEO of a PE-backed software company who is invisible in the AI conversation is conceding narrative ground to competitors who are not. Push PortCo CEOs onto panels, into ecosystems, into the conversation. It is talent acquisition, partnership development, and customer trust-building, bundled.

A 12-month operating calendar for the VCP

Horizon	Actions	Primary owner
Quarter 1	Stand up Engine 2 team with separate roadmap and direct CEO sponsorship; ship internal dogfooding of 2–3 agentic tools; redesign sales comp to favor AI SKUs.	CEO, with sponsor
Quarter 2	Publish three-layer KPI dashboard to the board; launch the platform-level AI advisory forum; design outcome-based / fractional-FTE pricing architecture.	CEO + operating partner
Quarter 3	Ship first AI-native product to GA; report measured productivity gains across the SDLC and back office; close the first ecosystem partnership or agentic tuck-in.	CEO + CTO, with VCP sponsorship
Quarter 4	Engine 2 ARR contributing material mix; demonstrate NRR uplift from AI-led expansion plays; reset next 12-month plan with sponsor; begin shaping the exit narrative around agentic ARR mix.	CEO + sponsor

SUMMARY: THE PLAYBOOK IS NOT NOVEL – THE EXECUTION IS

The survival pattern of every software cycle holds in this one too. Stay close to the customer. Operate at scale. Reinvest the savings into the new platform. What is different is the compression. Agentic AI is landing in production within quarters, not the decade cloud required. The switching-cost moats that funded two decades of PE software value creation are narrowing in real time, and the valuation compression that has accompanied this transition, while familiar in magnitude, is not familiar in speed.

For value creation teams, that compression has three consequences:

- **The investment thesis has to be rewritten at entry, not at exit.** AI-native readiness belongs on the first page of the IC memo, not in the technology appendix.

- **The VCP has to be redesigned around a two-speed operating model,** funded by deliberate margin compression and measured on a three-layer KPI dashboard.

- **Portfolio-level leverage is the sponsor's widest edge.** Platform advisory, shared GCC, programmatic bolt-ons, CEO bench, shared infrastructure. This is what public-market incumbents cannot replicate at the sponsor's speed.

What changes in the underwriting model

- **Revenue build.** Split ARR into legacy, AI-led, and agentic-outcomes tranches. Model each separately with its own growth rate, retention, and gross margin. The exit multiple will be a weighted function of the mix.

- **Margin trajectory.** Compression in the early hold (100–200 bps) followed by expansion (200-400 bps) as the AI factory matures. A flat trajectory is the wrong answer; it usually signals under-investment in Engine 2.

- **NRR as the primary EV driver.** Build the VCP around defending and expanding NRR. Sensitize the model on NRR, not new-logo bookings.

- **Exit comps and narrative.** The next buyer will separate legacy from AI-native ARR mix. Run a dual-track exit narrative from mid-hold onward.

The operating partners who program that trajectory deliberately rather than discovering it at exit readiness will define the next vintage's outperformers. The playbook is not novel. The execution is what will separate the next vintage from the repriced.

Copyright © L&T Technology Services