

VALUE VERTEX

*Scaling IT Services and Software Portfolios
to Peak Enterprise Value with AI*

THE PRIVATE EQUITY HANDBOOK FOR AI-DRIVEN VALUE CREATION

Contents

The Value Vertex — Introduction

Part I: The New Value-Creation Equation

Part II: Finding the Value

Part III: Software Portfolios

Part IV: IT Services Portfolios

Part V: The AI Workforce

Part VI: Governed Autonomy

Part VII: The 100-Day Value Creation Plan

Part VIII: The AI Value Office

Part IX: The PE Operating Model

Part X: The Exit

The Value Vertex Scorecard

The Value Vertex Playbook

Final Thought

The Value Vertex

Private equity has always been a business of transformation. Acquire a company. Identify the value-creation opportunities. Deploy capital and operational expertise. Accelerate growth. Expand margins. Build a stronger business. Exit at a higher enterprise value. That sequence has defined the industry for four decades, and it still does.

What has changed is the mechanism by which each of those steps gets executed. The traditional value-creation playbook was built around a world in which scaling a business generally required scaling its people, its infrastructure, and its operating complexity in rough proportion to its revenue. Grow the top line and the org chart grows with it. Add customers and you add implementation teams, support desks, and account managers. That relationship between growth and headcount has been, for most of the industry's history, treated as a law of physics rather than a design choice.

Artificial intelligence changes that equation. For software companies, AI can dramatically increase the velocity of product development, customer acquisition, implementation, support, and innovation — compressing timelines that used to take quarters into weeks. For IT services companies, AI can fundamentally change the relationship between revenue and headcount, moving the business from a labor-intensive model toward an AI-leveraged, and ultimately outcome-based, model. For both categories of asset, AI creates an opportunity to change not just how efficiently the company operates today, but what kind of company it can become by the time it changes hands again.

This is the Value Vertex — the point at which multiple value-creation vectors converge.

Growth. Margin. Capacity. Scalability. Differentiation. Risk. Exit readiness. AI can influence every one of these levers simultaneously, which is precisely what makes it different from prior generations of enterprise technology. Cloud migration made infrastructure more elastic. Automation software made discrete tasks faster. AI, deployed correctly, can touch the P&L, the balance sheet, the org chart, and the competitive position of a business at the same time — but only if it is deployed against the economics of the business rather than against a generic technology roadmap.

The objective of this handbook is therefore not “become an AI company.” That framing leads portfolio companies toward pilots, proofs of concept, and press releases that never touch EBITDA. The objective is: build a materially more valuable company because of AI. Everything that follows in this handbook is organized around that single distinction.

PART I

The New Value-Creation Equation

Why AI Changes the PE Playbook

CHAPTER 1

The fundamental equation of private equity has not changed: Enterprise Value = EBITDA × Multiple. What has changed is the number of ways AI can move both sides of that equation at once, and the speed at which it can do so.

$$\text{Enterprise Value} = \text{EBITDA} \times \text{Multiple}$$

How AI Moves EBITDA

On the earnings side, AI is a direct lever on the components that roll up into EBITDA. It can drive revenue growth by shortening sales cycles and improving win rates. It can expand gross margin by reducing the cost of delivering each unit of product or service. It can lift labor productivity, improve operating leverage, and increase utilization across delivery teams. It can lower cost-to-serve, reduce customer acquisition cost, accelerate product development, and reduce the ongoing cost of support. None of these are exotic technology outcomes — they are the same levers operating partners have always pulled, now with a materially more powerful tool.

- Revenue growth through faster, more targeted sales and marketing motions
- Gross-margin expansion through lower cost-to-serve
- Labor productivity gains across engineering, delivery, and support functions
- Improved operating leverage and utilization
- Reduced customer acquisition cost
- Faster product development cycles
- Lower support and service-desk costs

How AI Moves the Multiple

On the valuation side, the effect is more strategic and, for many sponsors, more consequential. Buyers pay premiums for businesses that look more scalable, more recurring, more differentiated, and less dependent on hard-to-retain labor. AI can influence scalability, recurring revenue mix, differentiation versus competitors, customer retention, defensible intellectual property, growth durability, operational resilience, labor dependency, and overall strategic attractiveness to the next buyer.

- Scalability of the operating model
- Recurring-revenue durability
- Competitive differentiation and defensibility
- Customer retention and expansion

- Proprietary IP and data assets
- Operational resilience
- Reduced dependency on hard-to-scale labor
- Strategic attractiveness to the next buyer

This is why AI should not sit inside the IT budget as one more technology initiative to be reviewed once a year alongside ERP upgrades and cybersecurity spend. It belongs inside the value-creation plan, reviewed with the same rigor as pricing strategy, M&A, or margin-improvement programs — because, done correctly, it will move the same numbers those programs are meant to move.

From AI Strategy to Value Strategy

CHAPTER 2

Most companies begin their AI journey by asking technology questions: Which model should we use? Should we deploy Copilot? Should we build agents? Which AI platform should we buy? Should we establish an AI center of excellence? These are reasonable questions for a CTO to ask, but they are the wrong starting point for a value-creation plan, because they start from the technology and work backward toward a business case that may or may not exist.

The PE question is different, and it is deceptively simple: where can AI create the greatest incremental enterprise value during the hold period? That single reframing changes the entire sequence of decisions that follow. Instead of beginning with a technology and searching for a use case, the sponsor begins with the investment thesis and searches for the AI lever that advances it.

Thesis-First Examples

If the thesis is to increase EBITDA by \$15 million over the hold period, the question becomes: where, specifically, can AI contribute to that \$15 million — in gross margin, in SG&A, in working capital efficiency? If the thesis is to grow revenue from \$100 million to \$150 million, the question becomes: where can AI increase growth velocity, whether through faster pipeline conversion, expanded account coverage, or a compressed sales cycle? If the thesis is to build a scalable services platform, the question becomes: how can AI decouple revenue growth from headcount growth, so the company can take on more work without adding cost in lockstep? If the thesis is to create a differentiated SaaS asset, the question becomes: how can AI change the product's underlying economics and its competitive position, not just its feature list?

AI becomes a lever against the thesis — not a parallel initiative running alongside it.

This distinction matters most at budget time. Initiatives that map cleanly to the thesis get funded, staffed, and measured against the same milestones as every other value-creation lever. Initiatives that cannot be traced back to the thesis tend to survive as pilots indefinitely, consuming management attention without ever appearing in a board deck as a dollar figure.

The Value Vertex Framework

CHAPTER 3

The Value Vertex framework evaluates every AI initiative across seven dimensions. No initiative needs to score well on all seven — but the strongest initiatives, the ones worth prioritizing first, tend to sit at the intersection of several vectors simultaneously rather than delivering a single, isolated benefit.

The Seven Vectors

- Growth — Can AI generate more revenue, whether through faster sales cycles, better targeting, or new product lines?
- Margin — Can AI reduce the cost of delivering revenue, in COGS or in operating expense?
- Capacity — Can the company produce more output without proportionally increasing headcount?
- Velocity — Can the company move faster across product releases, sales cycles, implementations, customer resolution, or decision-making?
- Differentiation — Can AI create capabilities that competitors cannot easily replicate?
- Risk — Can AI reduce operational, security, compliance, or execution risk?
- Exit — Will these improvements make the company demonstrably more attractive to the next buyer?

Consider an AI-enabled implementation platform for a services business. On its own, faster implementation sounds like an operations improvement. But trace the ripple effects: it reduces implementation cost (margin), it accelerates onboarding (velocity), it increases implementation capacity without adding headcount (capacity), it improves the customer experience during the highest-risk phase of the relationship (retention), and it reduces churn that would otherwise show up as lost revenue eighteen months later (growth and, ultimately, multiple). That combination — one initiative touching five or six vectors at once — is what this handbook calls a Value Vertex initiative, and it is the pattern operating partners should be actively hunting for during diligence and the first hundred days.

PART II

Finding the Value

Build the Portfolio Company Value Map

CHAPTER 4

Before deploying any AI initiative, the operating team needs a clear, quantified map of the company's underlying economics. This is not a new discipline for private equity — it is the same diligence rigor applied at acquisition, refreshed and pointed specifically at where AI can move the needle. Skipping this step is the single most common reason AI programs in portfolio companies fail to produce measurable results: teams start building before they know what “better” actually means in dollar terms.

Revenue Metrics

- Revenue and ARR
- Growth rate
- Retention and net revenue retention
- Average contract value
- Customer acquisition cost

Cost Metrics

- Delivery cost
- Engineering cost
- Support cost
- Sales and marketing spend
- G&A
- Infrastructure and cloud spend

Capacity Metrics

- Headcount by function
- Revenue per employee
- Gross profit per employee
- Customers per employee

Velocity Metrics

- Product release frequency
- Sales cycle length
- Implementation duration

- Support resolution time
- Decision-cycle time

Once these figures are assembled for every business unit, the picture that emerges is rarely uniform. One unit may be capacity-constrained; another may be margin-constrained; a third may be perfectly healthy on unit economics but bleeding customers due to slow implementation. The value map's purpose is to locate exactly where the company is constrained today, because that constraint — not the most exciting technology — is where the first AI dollar should go.

Follow the Work

CHAPTER 5

The most valuable AI opportunities are rarely visible from an org chart or a P&L. They are hidden inside workflows — the actual sequence of steps a person or team follows to move a piece of work from trigger to completion. Finding them requires literally mapping the work, process by process, the way an operations consultant would, rather than starting from a list of available AI tools and looking for somewhere to apply them.

Questions to Ask of Every Major Process

- What triggers the process?
- What information is required to complete it?
- Who performs the work, and at what seniority or cost?
- What decisions are made along the way?
- What systems are involved, and how much manual handoff exists between them?
- What output is produced?
- How long does the process take end to end?
- What percentage of the work requires rework or correction?
- Where are the bottlenecks?
- What happens to this process when volume doubles?

That last question deserves special attention. Many processes work acceptably well at current volume purely because people compensate for weak design with effort and tenure. The moment volume doubles — through organic growth, an add-on acquisition, or a new product line — those same processes break, and the company either hires aggressively to keep pace or begins to miss commitments. AI's greatest structural value is often not in making today's process faster, but in making it capable of absorbing volume growth without a proportional cost increase.

Work → Capacity → Cost → Revenue → EBITDA

AI becomes genuinely powerful, in the value-creation sense, precisely when it changes the shape of that chain — when doubling the work no longer requires doubling the people.

The AI Opportunity Matrix

CHAPTER 6

Once opportunities have been surfaced through the value map and the workflow audit, they need to be ranked with the same discipline applied to any capital allocation decision. The Value Vertex approach scores each opportunity across four dimensions.

- Economic Impact — how much value could this create, expressed in dollars against EBITDA, revenue, or capacity?
- Time to Value — how quickly can the benefit be realized and measured?
- Feasibility — can the technology, data, and team actually deliver this, given current maturity?
- Strategic Importance — does the initiative improve the company's long-term competitive position, independent of its near-term financial return?

$$\text{AI Value Score} = \text{Impact} \times \text{Speed} \times \text{Feasibility} \times \text{Strategic Importance}$$

The multiplicative structure of the score is deliberate: an initiative that scores a zero on feasibility should not survive simply because its theoretical impact is enormous. This keeps the prioritization exercise honest and prevents the most common failure mode in enterprise AI programs, which is chasing large but unachievable opportunities while ignoring smaller, immediately fundable ones. The highest-scoring initiatives, taken together, become the portfolio company's AI value-creation agenda for the hold period — a living list, revisited quarterly, not a one-time exercise filed away after the first offsite.

PART III

Software Portfolios

Building the AI-Native Software Company

CHAPTER 7

Software businesses have an unusual opportunity relative to other portfolio companies: AI can touch essentially every stage of the software value chain, from the earliest product discovery work through to the customer's day-to-day support experience. Sponsors who treat this as a single “adopt AI in engineering” initiative are leaving most of the opportunity on the table.

Product

- Analyze customer feedback at scale
- Identify product opportunities from usage and support data
- Generate specifications and requirements drafts
- Detect product gaps against competitors
- Predict feature adoption before a release ships

Engineering

- Generate code from specifications
- Review code for quality and security issues
- Create automated tests
- Troubleshoot and triage defects
- Modernize legacy systems and reduce technical debt
- Maintain documentation as the codebase evolves

Quality Assurance

- Generate test scenarios from requirements
- Analyze failures to identify root cause faster
- Prioritize regression testing based on risk
- Detect quality risks before release

DevOps

- Investigate incidents in real time
- Monitor systems for anomalies
- Recommend remediation steps

- Automate routine operational tasks

Customer Success

- Identify churn signals earlier in the lifecycle
- Generate account intelligence for CSMs
- Recommend expansion opportunities
- Automate routine customer communications

Support

- Resolve common issues automatically
- Diagnose problems faster
- Retrieve knowledge instantly across a large knowledge base
- Create and route tickets to the right owner

The opportunity here is not simply better productivity within each of these functions taken separately. Stacked together, across product, engineering, QA, DevOps, customer success, and support, the opportunity is a fundamentally different software operating model — one where the cost to build, ship, and support the product no longer scales linearly with the size of the customer base.

The Software Development Multiplier

CHAPTER 8

Software companies have historically scaled by adding engineers. Revenue grows, the roadmap expands, and headcount follows. AI introduces a second scaling mechanism entirely: engineering capacity without proportional headcount growth.

Consider a hypothetical hundred-person engineering organization that produces \$50 million of annual product value. If AI increases effective engineering capacity by 30 percent, management now has a genuine strategic choice to make, not merely a cost-reduction opportunity to bank. It can reduce development cost and drop the savings to the bottom line. It can ship more features and accelerate competitive differentiation. It can accelerate roadmap delivery to hit a customer commitment sooner. It can redirect the freed-up capacity toward entering new markets. It can spend the capacity on quality and reduce the rate of production incidents. Or it can use it to pay down accumulated technical debt that has been slowing every subsequent release.

- Reduce development cost and improve margin
- Ship more features and widen the competitive gap
- Accelerate roadmap delivery against commitments
- Enter new markets with freed-up engineering capacity
- Improve quality and reduce production incidents
- Reduce accumulated technical debt

Productivity is not the value. Productivity is the capacity from which value can be created.

The right answer among those choices depends entirely on the investment thesis established in Part I. A sponsor pursuing a growth thesis should generally push the capacity toward roadmap and market expansion; a sponsor pursuing a margin-improvement thesis ahead of a near-term exit should generally push it toward cost. What must not happen is for the capacity to simply evaporate into faster status quo — measured productivity gains that never convert into a decision and, therefore, never convert into value.

From SaaS to AI-Native Software

CHAPTER 9

The most consequential transformation available to a software portfolio company may not occur inside engineering at all — it may occur at the level of the product itself. Traditional SaaS software helps customers perform work: it organizes their data, structures their workflow, and gives them tools to execute faster. AI-native software does more of the work itself, on the customer's behalf.

That distinction sounds subtle, but it creates entirely new product economics. A customer historically wanted “a system where my employee manages the workflow.” Increasingly, that same customer wants “a system that manages the workflow for me.” The shift moves the competitive battlefield away from feature checklists and toward outcomes delivered — a fundamentally different basis for pricing, packaging, and retention.

System of Record + System of Intelligence + System of Action

The software companies best positioned to win the next several years will increasingly combine all three layers: a system of record that captures and organizes the customer's data, a system of intelligence that reasons over that data, and a system of action that actually executes work rather than merely recommending it. Portfolio companies still operating purely as a system of record are the most exposed to disruption from AI-native competitors, and diligence on any software asset should explicitly assess how far along this spectrum the product currently sits.

PART IV

IT Services Portfolios

The IT Services Value Equation

CHAPTER 10

Traditional IT services economics are relatively simple to model: revenue is a function of people, their utilization, and their bill rate.

$$\text{Revenue} = \text{People} \times \text{Utilization} \times \text{Rate}$$

AI introduces a new variable into that equation — AI-leveraged capacity — that sits between people and utilization and changes what a given headcount is capable of producing.

$$\text{Revenue} = \text{People} \times \text{AI-Leveraged Capacity} \times \text{Utilization} \times \text{Rate}$$

This single addition creates the possibility of decoupling revenue growth from headcount growth, which is arguably the single most valuable structural shift available to a services business under private equity ownership. Consider a \$200 million services business in which AI increases productive capacity by 25 percent. Management has effectively created the equivalent of \$50 million of additional delivery capability without adding a proportional number of employees.

That new capacity is not automatically value — it is optionality that must be deliberately allocated. It can be used to grow revenue by taking on more client work with the existing base. It can be used to increase margins by delivering the same revenue with a lighter cost structure. It can be used to reduce prices strategically in order to win share. It can be used to increase utilization of the existing bench. It can be used to accelerate delivery timelines as a competitive differentiator. Or it can be used to create entirely new service offerings that were previously uneconomical to staff. The strategic decision about which of these paths to pursue is, again, dictated by the investment thesis — but the decision itself is now available to management in a way it simply was not five years ago.

From Labor Arbitrage to AI Arbitrage

CHAPTER 11

The IT services industry has historically competed primarily through labor economics: onshore versus nearshore versus offshore delivery, and access to lower-cost talent pools. AI introduces a new and more durable form of competitive advantage — intelligence arbitrage. The winning provider going forward may not be the one with access to the cheapest people, but the one that can deliver the same business outcome with dramatically less human effort of any cost tier.

The Progression

This shift tends to move through a recognizable progression as a services company matures its AI capability.

- Labor Arbitrage — competing on lower-cost people
- Technology Arbitrage — competing on better tools
- AI-Leveraged Delivery — people working alongside AI
- AI-Native Delivery — AI performs substantial portions of the workflow directly
- Outcome-Based Services — customers pay for results rather than hours or headcount

That final transition, from labor-based billing to outcome-based pricing, may represent the single largest structural opportunity available to the services industry over the next several years. It changes the unit of value the customer is buying, and it changes the ceiling on the margin the provider can capture, because the provider is no longer capped by the number of billable hours it can staff.

The End of the Headcount Growth Model

CHAPTER 12

For decades, services companies have operated under a structural equation that every operating partner in the industry has internalized: more revenue requires more people. AI creates the possibility of an entirely different equation: more revenue requires more AI-enabled capacity, not more people.

Consider a services company that previously needed one hundred people to deliver \$10 million of revenue. Under an AI-leveraged delivery model, that same workforce might eventually deliver \$15 million of revenue without adding headcount. It is worth being explicit that this does not automatically mean eliminating a third of the workforce — and framing it that way is usually the wrong strategic choice for a company mid-hold-period. The company may instead redeploy that freed capacity toward increasing revenue, improving customer experience, entering new markets, reducing prices to win share, or increasing margins outright.

The best PE sponsors will recognize that capacity creation can be more valuable than cost reduction.

This is a genuinely important strategic distinction. A sponsor who treats every productivity gain as a headcount reduction opportunity will capture a one-time margin improvement and then plateau, having signaled to the market that the business is a cost story rather than a growth story. A sponsor who treats the same productivity gain as capacity to be redeployed toward growth builds a business that is both more valuable at exit and more resilient during the hold period.

AI-Native Managed Services

CHAPTER 13

The most interesting structural opportunity for a services portfolio company may be replacing labor-based staffing contracts with AI-enabled managed outcomes — changing not just how the work gets done, but what the company is actually selling.

Reframing the Offer

Instead of “we provide twenty ServiceNow engineers,” the company offers “we operate your IT service management function.” Instead of “we provide fifty customer support agents,” the company offers “we manage your customer resolution operation.” Instead of “we provide a development team,” the company offers “we operate your software delivery capability.”

Each of these reframings changes five things at once: pricing structure, margin profile, the nature of the customer relationship, the scalability of the engagement, and ultimately the valuation the market is willing to assign to the business. A staffing-model services provider is priced like a staffing company, with revenue capped by headcount and margins capped by wage inflation. An outcomes-model provider begins to look, and trade, more like a technology platform — because the customer is now buying a managed result rather than a body count, and the provider's economics scale with the efficiency of its delivery model rather than with the number of people it can hire.

PART V

The AI Workforce

Humans + AI Agents

CHAPTER 14

The future portfolio company will not be entirely human, nor will it be entirely autonomous — it will be a hybrid organization, and the design of that hybrid is itself a source of competitive advantage.

What Humans Own

- Strategy
- Accountability
- Relationships
- Judgment
- Exceptions
- High-impact decisions

What AI Agents Perform

- Analysis
- Research
- Coordination
- Documentation
- Execution
- Monitoring
- Repetitive decisions

This division of labor gives rise to a new organizational unit worth designing deliberately: the AI POD, which combines human leadership with specialized AI workers organized around a single outcome, rather than around a traditional functional department.

- Software Delivery POD — product leader + engineers + AI development workers + AI testing + AI operations
- Customer Operations POD — customer leader + support specialists + AI support workers + AI knowledge agents
- Finance POD — controller + finance specialists + AI reconciliation, reporting, and forecasting agents
- Sales POD — sales leader + account executives + AI research, prospecting, and proposal agents

The unit of scale becomes the POD, not the individual employee.

This has direct implications for how a portfolio company plans headcount during the hold period. Instead of forecasting individual roles to add as revenue grows, operating teams should be forecasting PODs to stand up — each with a defined outcome, a human accountable owner, and a specified mix of AI capability supporting it.

Redesign Jobs Around Outcomes

CHAPTER 15

The wrong question to ask when introducing AI into an organization is “which jobs can AI replace?” That question invites defensiveness from the organization, produces incomplete answers, and typically results in shallow automation of the easiest, least valuable tasks. The right question is “what outcome does this role own?”

Once a role's outcome is clearly defined, it can be decomposed into its constituent parts: the decisions it makes, the tasks it performs, the workflows it participates in, the communications it is responsible for, the analysis it conducts, the execution it carries out, and the exceptions it handles. Each of those components can then be classified as best suited to a human, an AI system, or some combination of the two.

- Human
- AI
- Human + AI

The objective is not maximum automation. It is maximum economic leverage.

This distinction protects against the most common overcorrection in AI-driven redesign, which is treating every task that can technically be automated as a task that should be automated regardless of the judgment, relationship, or accountability it requires. A role redesigned for maximum economic leverage will often look like a smaller number of people spending a larger share of their time on the decisions and relationships that actually require human judgment, supported by AI handling the volume work around them.

PART VI

Governed Autonomy

The Autonomy Problem

CHAPTER 16

AI dramatically increases execution capability, and that increase in capability creates a corresponding increase in risk that must be actively managed rather than assumed away. An AI assistant that drafts a document can simply be wrong. An AI agent that changes production infrastructure can cause an operational incident. An AI agent empowered to approve a customer refund can create direct financial exposure.

As AI systems move from generating information to taking action on behalf of the business, governance stops being a compliance afterthought and becomes a foundational design requirement. The question every operating team must answer shifts from “can AI do this?” to “should AI be allowed to do this, under what conditions, and with what oversight?”

AI Governance as a Value-Creation Capability

CHAPTER 17

It is tempting to treat governance purely as a defensive, compliance-oriented cost center. That framing understates its actual value. Properly designed, governance is what enables an organization to safely extend more autonomy to its AI systems — and the amount of work a company can confidently delegate to AI is, in practice, gated by the strength of the controls surrounding it, not by the raw capability of the underlying models.

What a Governed AI Worker Requires

- Identity — a clear, auditable record of which system is acting
- Authorization — explicit permission scoped to specific actions
- Defined policies — rules that constrain what the system can and cannot do
- Action boundaries — hard limits on the scope and magnitude of any single action
- Human approval rules — checkpoints for actions above a defined risk threshold
- Auditability — a complete, reviewable record of every action taken
- Monitoring — ongoing visibility into system behavior
- Escalation — a clear path to a human when the system encounters an exception

Maximum Safe Autonomy.

That phrase — maximum safe autonomy, rather than maximum autonomy — is the correct target. The more confidence management has in the controls surrounding a given AI system, the more work it can responsibly delegate to it, and the faster the capacity gains described in earlier chapters can actually be realized in production rather than left stalled in a pilot phase out of an abundance of caution.

The AI Control Plane

CHAPTER 18

As a portfolio company moves from a handful of AI pilots to dozens or hundreds of AI workers embedded across its operations, it needs a control layer capable of answering a consistent set of questions across every one of those systems, on demand.

- Who — or what — is acting?
- On whose behalf is it acting?
- What is it trying to accomplish?
- What is it allowed to do?
- What did it actually do?
- Did that action comply with policy?
- What evidence exists to demonstrate that compliance?

This becomes particularly important for PE-backed companies specifically, because autonomous execution introduces a genuinely new category of operational risk — one that did not exist, at this scale, in the prior generation of enterprise software. A misconfigured integration could previously cause a data error that a person would eventually catch. A misconfigured AI agent with write access to production systems can take thousands of incorrect actions before anyone notices. The control plane is what provides management, the board, and eventually a buyer, with the confidence required to scale autonomy rather than being forced to keep it artificially small out of caution.

PART VII

The 100-Day Value Creation Plan

Days 1–30: Diagnose

CHAPTER 19

The first thirty days after committing to an AI-driven value-creation program should be spent establishing a rigorous, quantified baseline — not deploying technology. Every subsequent claim of impact will be measured against the numbers captured in this window, so the discipline applied here determines the credibility of everything that follows.

- Revenue
- EBITDA
- Gross margin
- Headcount
- Revenue per employee
- Utilization
- Customer retention
- Support cost
- Engineering throughput
- Delivery cycle time

Alongside the baseline, the team should map the highest-value workflows using the method described in Chapter 5, and use the AI Opportunity Matrix from Chapter 6 to identify and rank the top twenty candidate initiatives. The output of this phase is not a technology decision — it is a prioritized, quantified list ready to be tested.

Days 31–60: Prove

CHAPTER 20

From the top twenty candidates identified in the diagnostic phase, select three to five initiatives to actually build and deploy. Each selected initiative must have a named executive owner, a documented financial baseline, a specific target outcome, a technology owner responsible for delivery, a risk owner responsible for governance, and a realistic implementation timeline.

- Executive owner accountable for the business outcome
- Financial baseline captured before deployment
- Target outcome, expressed in measurable terms
- Technology owner responsible for delivery
- Risk owner responsible for governance and controls
- Implementation timeline with clear milestones

Critically, these initiatives must be deployed into real, live workflows — not staged as demonstrations for a leadership offsite. The objective of this phase is not to optimize for an impressive demo; it is to optimize for measurable economic impact that can be defended in front of a board or an investment committee.

Days 61–100: Scale

CHAPTER 21

With three to five initiatives proven in real workflows, the final phase of the first hundred days is about scaling what worked and, just as importantly, killing what did not. Scaling successful initiatives means building the reusable infrastructure that lets the same capability be deployed across additional teams, business units, or customer segments without rebuilding it from scratch each time.

- Agents
- Integrations
- Knowledge bases
- Policies
- Workflows
- Evaluation frameworks

Pilots that failed to produce measurable value should be killed decisively rather than allowed to linger in a permanent “experimental” state that quietly drains engineering time and management attention. By the end of day one hundred, the portfolio company should be able to point to validated AI use cases with measured financial impact, a repeatable deployment model that does not depend on a single hero engineer, an operating AI governance framework, and a prioritized twelve-to-twenty-four-month roadmap that ties directly back to the investment thesis.

PART VIII

The AI Value Office

From Transformation Office to Value Office

CHAPTER 22

Traditional transformation offices exist to manage projects: they track milestones, status, and delivery risk. The AI Value Office that this handbook recommends portfolio companies stand up is designed to manage economic outcomes instead — its dashboard is built around dollars and business metrics, not project status.

Financial

- EBITDA impact
- Revenue impact
- Cost reduction
- Margin expansion

Productivity

- Output per employee
- Capacity created
- Automation rate

Velocity

- Product cycle time
- Sales cycle
- Implementation time
- Resolution time

Customer

- Retention
- Expansion
- Net Promoter Score
- Customer effort

AI System Health

- Agent success rate
- Intervention rate

- Error rate
- Cost per transaction
- Policy violations

The most important metric is Realized Value — not estimated value, not potential value, not AI adoption.

This last distinction is the one that separates portfolio companies that genuinely capture AI-driven value from those that generate an impressive-looking slide of “AI adoption” metrics that never translate into a board-reportable financial result. Adoption rate is a leading indicator at best; it is not a substitute for realized value, and the AI Value Office should be judged, and should judge itself, on the latter.

The AI ROI Waterfall

CHAPTER 23

To keep every AI business case grounded in reality rather than in the initial, most optimistic estimate, each initiative should be tracked as it moves through five progressively more conservative stages.

- Gross Opportunity — the theoretical maximum value if the problem were solved perfectly
- Addressable Opportunity — the portion of that value AI can actually influence
- Implementable Opportunity — the portion realistically achievable given current data, systems, and team maturity
- Realized Opportunity — the improvement actually measured in production
- Net Value — realized value minus the full cost of AI, software, implementation, integration, governance, change management, and ongoing human supervision

Reporting only the top of this waterfall — the gross opportunity — is how enterprise AI programs earn a reputation for overpromising and underdelivering. Reporting the full waterfall, with net value as the number that actually matters to the investment committee, is what keeps the AI Value Office credible over multiple board cycles.

PART IX

The PE Operating Model

AI Due Diligence

CHAPTER 24

AI should increasingly be treated as a formal workstream within technology and operational due diligence, evaluated with the same structure applied to any other diligence category, and assessed for both upside and exposure before capital is committed.

Market

- Customer AI exposure
- AI-driven disruption risk in the category
- Competitive threats from AI-native entrants

Product

- AI roadmap maturity
- Proprietary data assets
- AI-driven differentiation versus competitors

Technology

- Architecture readiness for AI integration
- Technical debt
- Engineering team maturity

Operations

- Labor intensity of core processes
- Workflow automation potential
- Specific AI agent opportunities

Workforce

- Critical talent dependencies
- Skill scarcity risk
- Organizational AI readiness

Risk

- Data quality and rights

- Cybersecurity posture
- Regulatory and compliance exposure
- Existing governance maturity

The objective: identify both AI upside and AI exposure before investment.

The First 30 Days After Close

CHAPTER 25

Immediately after close, the operating team should establish the AI baseline described in Chapter 19, and begin working through the same set of strategic questions that framed the Value Vertex itself: where can AI create EBITDA, where can it accelerate growth, where can it create capacity, where can it reduce customer friction, where can it create differentiation, and where can it reduce risk?

- Where can AI create EBITDA?
- Where can AI accelerate growth?
- Where can AI create capacity?
- Where can AI reduce customer friction?
- Where can AI create differentiation?
- Where can AI reduce risk?

Every opportunity surfaced by those questions should then be prioritized according to the hold period. A sponsor should favor opportunities that can demonstrate measurable impact within the investment horizon, while still building strategic value that will matter to the eventual buyer — a balance that requires discipline, since the most exciting long-horizon opportunities are not always the ones that move the numbers a buyer will diligence three or four years from now.

The Hold-Period Roadmap

CHAPTER 26

Across a typical hold period, AI value creation follows a recognizable arc, moving from quick, provable wins toward genuine business-model transformation.

0–100 Days: PROVE

Find quick wins. Establish the baseline, run the diagnostic, and prove three to five initiatives in real workflows.

3–6 Months: SCALE

Industrialize the workflows that were proven in the first phase, extending them across additional teams and business units.

6–12 Months: REDESIGN

Move beyond point solutions and begin changing the underlying operating model — the org design, the job architecture, and the workflow structure themselves.

12–24 Months: DIFFERENTIATE

Build AI-native products and services that competitors without the same capability cannot easily replicate.

Exit: MONETIZE

Translate everything built over the prior phases into the equity story that will be told to the next buyer.

PART X

The Exit

Building the AI Equity Story

CHAPTER 27

At exit, the story a sponsor tells about AI should never be “we implemented AI.” That framing describes an activity, not an outcome, and sophisticated buyers will discount it accordingly. The correct story is: “we changed the economics of the company.”

For a Software Company

- Engineering throughput increased 40 percent
- Release cycles decreased 35 percent
- Support cost declined 25 percent
- Customer retention improved measurably
- AI created new, incremental product revenue

For an IT Services Company

- Revenue per employee increased
- Delivery margins expanded
- Project cycle times declined
- Utilization improved
- AI-enabled services generated new revenue streams
- Headcount growth decoupled from revenue growth

Each of those bullet points is a specific, defensible, diligence-ready claim — not a vague assertion of AI adoption. That specificity is what turns an AI transformation narrative into an actual exit story that supports a premium multiple.

The AI-Native Portfolio Company

CHAPTER 28

The strongest portfolio companies emerging from this era will share a common set of characteristics: AI-enabled revenue generation, AI-enabled delivery, AI-enabled operations, AI-enabled management, governed AI autonomy, and genuinely AI-native products — all of it converging into measurably superior economics rather than remaining a collection of disconnected pilots.

- AI-enabled revenue
- AI-enabled delivery
- AI-enabled operations
- AI-enabled management
- Governed AI autonomy
- AI-native products

The objective is not to become the company with the most AI. It is to become the company with the best economics because of AI.

The Value Vertex Scorecard

Score each dimension from one to five to establish where a portfolio company currently stands, and to track its progress toward the Value Vertex over the course of the hold period.

Dimension	Guiding Question	Score (1–5)
Strategy	Is AI explicitly tied to the investment thesis?	
Growth	Is AI increasing revenue velocity?	
Margin	Is AI expanding gross or EBITDA margin?	
Capacity	Is output increasing without proportional headcount?	
Velocity	Are critical business cycles getting faster?	
Product	Is AI creating differentiated offerings?	
Operations	Are workflows being redesigned, not just accelerated?	
Autonomy	Can AI safely execute meaningful work?	
Governance	Is autonomous work controlled and auditable?	
Exit	Can AI impact be demonstrated to a future buyer?	

Interpreting the Score

- 1–2: Experimenting — isolated pilots with no clear financial link
- 3: Operationalizing — initiatives proven in production but not yet scaled
- 4: AI-enabled — AI measurably improving core economics across functions
- 5: AI-native — AI, operating model, economics, and strategy reinforcing one another

The objective is to reach the Value Vertex itself: the state in which AI, operating model, economics, and strategy reinforce one another rather than operating as separate initiatives competing for the same budget and attention.

The Value Vertex Playbook

Ten steps, applied in sequence, translate everything in this handbook into an operating program.

Step 1: Start with the investment thesis

Every AI initiative must trace back to a specific element of the thesis before it is funded.

Step 2: Map the economics

Build the portfolio company value map described in Chapter 4.

Step 3: Map the work

Follow the work through real processes, as described in Chapter 5.

Step 4: Identify AI leverage points

Score opportunities using the AI Opportunity Matrix in Chapter 6.

Step 5: Quantify the value

Attach a specific, defensible dollar figure to each candidate initiative.

Step 6: Prove value in 100 days

Follow the diagnose, prove, and scale sequence in Part VII.

Step 7: Scale successful workflows

Industrialize what worked; kill what didn't.

Step 8: Redesign the operating model

Move beyond point solutions to job, workflow, and org redesign.

Step 9: Govern autonomous work

Build the control plane described in Part VI before extending autonomy further.

Step 10: Build the exit story

Translate realized value into the specific, diligence-ready narrative described in Chapter 27.

Final Thought

Private equity has always been about finding the levers that make a company more valuable. AI introduces a new class of leverage. It can increase the output of people. It can reduce the cost of work. It can accelerate decision-making. It can compress cycle times. It can create new products. It can transform services economics. It can enable entirely new operating models.

But AI itself is not the destination. The destination is enterprise value.

The winning PE firms will stop asking “How should our portfolio companies adopt AI?” and start asking “How can AI change the economic trajectory of this company?”

The winning portfolio-company CEOs will ask what work should be performed by humans, what work should be performed by AI, and how the business should be redesigned around that new division of labor. And the winning operating partners will ask how to turn those changes into measurable EBITDA, growth, scalability, differentiation, and exit value.

That is the discipline of Value Vertex. Not AI for AI's sake. Not transformation theater. Not another technology program. A systematic approach to using AI to move a portfolio company toward peak enterprise value.

Find the Value. Accelerate the Value. Reach the Vertex.