

FLOAT

Firm Leakage Opportunity And Triage

An Agentic AI EBITDA Playbook for CFOs and PE Operating Partners

90-Minute Diagnostic to Find \$5.3+ Million of Hidden EBITDA In Mid-market \$50-500M Firms

Before you venture into the SEAS, make sure you can FLOAT.

**For Mid-Market CFOs and PE Operating Partners
Who Are Tired of Being Asked
"Where Is the EBITDA from All This AI?"**

LALIT KUMAR

B.Sc. (Maths, Physics, Computer Science), MCA, LL.B.

Author of the Strategic EBITDA Acceleration System (SEAS)

www.smartagenticsystems.com

Copyright

© 2026 Lalit Kumar. All rights reserved.

FLOAT (**F**irm **L**eakage **O**ppportunity **A**nd **T**riage), the FLOAT Diagnostic System, the EBITDA Leakage Framework, the Strategic EBITDA Acceleration System (SEAS), the SEAS Playbook, and the RAPID SEAS Playbook and SPEED Playbook are trademarks of Lalit Kumar.

No part of this book may be reproduced, distributed, transmitted, or stored in any form or by any means without the prior written consent of the author, except for brief quotations in critical reviews and certain other non-commercial uses permitted by copyright law.

Disclaimer: This book is provided for informational and strategic planning purposes only. It does not constitute legal, financial, tax, or professional advice. The financial projections, EBITDA improvement estimates, and performance metrics in this book are based on industry research, published case studies, and modelling assumptions. Actual results will vary based on implementation quality, organisational readiness, data quality, market conditions, and management commitment. Readers should consult qualified professionals before making business decisions based on the content herein.

AI Disclosure: This book was written by the author. All intellectual property — the SEAS framework, the FLOAT diagnostic methodology, the financial models, the canonical figures, the three-pocket leakage taxonomy, the gate criteria, and the 14-day execution playbook — originated with the author across two years of independent research and synthesis. It is not AI generated but was prepared with strictly regulated AI assistance for prose drafting, structural editing, and copy refinement only.

All third-party research citations (Bain & Company, McKinsey & Company, Boston Consulting Group, FTI Consulting, KPMG, Accenture, Gartner, RAND, NIST) are attributed to their respective authors and organisations and are used for educational purposes under fair use. The author does not claim affiliation with, endorsement by, or sponsorship from these organisations.

Evidence Quality Notice: Evidence quality for individual levers varies. The full SEAS Playbook contains source verification ratings for each lever; buyers should review

lever evidence ratings in the SEAS Playbook before including any lever in IC presentations.

First edition, 2026.

Why I Wrote FLOAT

I spent two years building the Strategic EBITDA Acceleration System (SEAS). More than eight hundred pages. Thirty-five agent specifications. Nine financial models. Twenty legal templates. A complete implementation system for capturing the operational EBITDA that mid-market companies leak into the \$50M to \$500M revenue band, every month, without noticing.

When I finished SEAS, I noticed a pattern that surprised me. The CFOs and PE operating partners I described it to did not need 800 pages first. They needed 90 minutes. They needed to know whether the leakage was real, how big it was for their specific firm, and which pocket to attack first. The 800 pages mattered, but only after the first 90 minutes were done. The diagnostic had to come before the playbook.

So I extracted the diagnostic. I made it stand on its own. I wrote it in a way that a sceptical CFO could finish in a single sitting and run on a Tuesday morning before lunch. I gave it a name — FLOAT, Firm Leakage Opportunity And Triage. And I priced it at the cost of a printer cartridge.

The math behind FLOAT is the same math that anchors SEAS. The canonical figures are the same: a \$150 million revenue mid-market firm, \$18 million of EBITDA, 12% margin, \$5.3 million of recoverable leakage hidden across three pockets that your existing P&L already documents. The arithmetic is not ambitious. It is conservative. The high-performer figures published by Bain, FTI Consulting, and KPMG are higher than what FLOAT and SEAS use as base cases, because I would rather understate the recovery and watch you exceed it than the alternative.

If after reading this book you reach for SEAS, that is the path the architecture is designed for. If you finish FLOAT, run the diagnostic, and decide to attack the pockets without the implementation system, that is also legitimate. The diagnostic is yours. The worksheet you produce in 90 minutes is yours. The decision about what comes next is yours.

What is not negotiable is the diagnostic itself. Every month you spend not running it is \$375,000 of margin that does not come back. That is the only sentence in this book I want you to remember.

— Lalit Kumar
June 2026

Table of Contents

Foreword:	The \$375,000 You Are Losing This Month.....	7
Chapter 1:	The \$5.3 Million Hiding in Your P&L.....	10
Chapter 2:	Why Three Out of Four AI Initiatives Fail.....	14
Chapter 2.5:	The 2026 Mid-Market Operating Reality.....	17
Chapter 3:	FLOAT — The Framework.....	19
Chapter 4:	Pocket 1 — The SG&A Efficiency Gap.....	21
Chapter 5:	Pocket 2 — Pricing and Contract Drift.....	24
Chapter 6:	Pocket 3 — Supply Chain and Vendor Concentration	27
Chapter 7:	Sizing the Prize — Your ROI Calculator.....	30
Bonus:	Walking Through a Filled Worksheet.....	32
Chapter 8:	The Three Stress Tests Every IC Will Run.....	35
Chapter 8.5:	Common Diagnostic Mistakes.....	37
Chapter 9:	From Diagnostic to Pilot in 14 Days.....	40
Bonus:	From Worksheet to Board Memo.....	42

Chapter 10: What Comes After FLOAT — The SEAS System.....	45
Bonus: Frequently Asked Questions.....	50
About the Author	53
Your Next Step	54
Appendix A: Worksheet	55
 The Strategic EBITDA Acceleration System	62
Execution Flow: diagnose (Ch 1-3) → fill in details (Ch 4-6) → measure (Ch 7) → defend (Ch 8) → execute (Ch 9) → scale (Ch 10)	

How to Use This Book

This book is structured as a single 90-minute working session, not a leisurely read. The chapters correspond to the stages of the diagnostic, in the order you will run them.

Chapters 1 and 2 set the context. Read these first to understand what you are looking for and why most companies miss it. Together, they are about 25 minutes of reading.

Chapter 3 introduces the FLOAT framework — the four-by-three diagnostic grid that produces the worksheet output. Read this carefully. Five minutes.

Chapters 4, 5, and 6 walk through the three pockets in turn. Each chapter takes about ten minutes to read and gives you the red-flag thresholds, the dollar quantification formula, and the candidate workflow for that pocket. After reading them, you are ready to run the diagnostic.

Chapter 7 is the ROI calculator. Use this when you have your three pocket sizes and need to pick the pilot.

Chapter 8 is the script for the IC stress tests you will face. Internalise this before you walk into the board meeting.

Chapter 9 is the 14-day playbook from completed diagnostic to pilot kickoff. Chapter 10 explains what SEAS is and why you might want it.

Appendix A is the printable worksheet. Print it before you start the diagnostic.

The fastest path through the book is to read Chapters 1 to 6 first, then sit down at a table with the trailing twelve-month P&L, the SG&A breakdown, and the vendor concentration report. Set a 90-minute timer. Run the diagnostic. Then read Chapters 7 to 10 with your worksheet in front of you to size and pilot what you have found.

If you take more than 90 minutes for the diagnostic itself, the diagnostic has drifted into a workshop. Stop. Compress what you have onto the worksheet. Move on.

Foreword: The \$375,000 You Are Losing This Month

If you run a \$150 million mid-market company, and you have not run a structured EBITDA leakage diagnostic in the last twelve months, the math is not in dispute.

This playbook uses a \$150 million revenue company as the worked example throughout, because it is the median of the SEAS target range (\$50M-\$500M). The methodology and dollar logic scale linearly across the full range — firms at the upper end typically identify \$8M-\$15M+ in the same diagnostic, while firms at the lower end identify \$1.5M-\$3M+. The \$5.3M figure used throughout is the median anchor, not a ceiling. Substitute your own revenue and the percentages hold.

Three percentage points of revenue leakage is the median in mid-market firms that have not specifically addressed it. Three percent of \$150 million is \$4.5 million a year. Divided by twelve months, that is \$375,000 a month. Divided by thirty days, that is \$12,500 a day. Chapter 1 produces a slightly higher figure of \$5.3 million by quantifying each leakage category bottom-up — the two numbers reconcile because the bottom-up waterfall captures pockets the conservative 3% benchmark rounds away. The Foreword math uses the conservative number; Chapter 1 onwards uses the canonical \$5.3M.

That is not a projection. That is not a consultant's pitch. That is what the trailing-twelve-month P&L of a normally-run mid-market company shows when you read it through the right lens, and it is the reason this book exists.

Every month you spend reading one more McKinsey article, scheduling one more vendor demo, or waiting for the next budget cycle, you are burning \$375,000 in unrecovered margin. That money is not coming back. It leaks out the bottom of the P&L and it is gone the moment the calendar turns. What you walk away with after 90 minutes is not a vague sense that AI might help. It is a single page with three EBITDA pockets sized in dollars against your own P&L, ranked by speed-to-cash, with the highest-ROI one named as your candidate pilot workflow. That single page is the artifact you take to your CFO, your CEO, or your IC. If you run nothing else from this book, you will have that page. The page is what changes the conversation from "we should look into AI" to "we are pursuing a \$1.2M annualised opportunity in AP automation, here is the workflow, here is the vendor shortlist, here is the timeline." That is the value of the diagnostic, independent of whether you ever buy SEAS.

This book exists to stop that haemorrhage within two weeks of you reading it. Not through a transformation programme. Not through a six-month consulting engagement. Through a 90-minute diagnostic that uses data your CFO already has on hand, identifies three quantifiable EBITDA pockets by lunch, and produces a named pilot workflow by Friday — followed by a fourteen-day path to vendor selection and pilot kickoff (Chapter 9).

The diagnostic is called FLOAT — Firm Leakage Opportunity And Triage. The output is a single page with three EBITDA pockets sized in dollars, not adjectives. The input is the trailing twelve-month P&L, an SG&A breakdown, and a vendor concentration report. If you cannot pull those three documents in fifteen minutes, your data infrastructure is the actual problem, and FLOAT will tell you that in the first thirty minutes — which is also useful information.

A short note on what this book is and is not.

It is not the Strategic EBITDA Acceleration System (SEAS) itself. SEAS is the 309-page master playbook plus 9 Excel models, 35 strategic agent specifications, 20+ legal templates, and the board-ready presentation decks that carry an organisation from the validated diagnostic through the 18-month deployment, the five go/no-go gates, the NIST-aligned governance framework, the vendor SLA templates, and the exit documentation. SEAS is the treatment plan. FLOAT is the diagnosis. You have paid FLOAT's price, or zero if a colleague forwarded you the file. You have not paid for SEAS, the implementation system, which is \$2,997. The relationship between the two is that of a free X-ray and a paid surgeon. The X-ray is genuinely useful even if you decide to consult a different surgeon, or no surgeon at all. The 90-minute diagnostic in this book is the same. It produces the diagnostic worksheet whether or not you go on to purchase the implementation system, and the worksheet is yours to keep, share with your CFO, present to your board, or file in a drawer. There is no upsell embedded in the diagnostic itself. The upsell, if it comes, is the moment you realise that running the diagnostic was the easy part.

You do not skip the diagnosis. You also do not stop there, because a diagnosis without a treatment plan is just an expensive piece of paper.

This book is the diagnosis. Read it once. Run it once. Then, if the diagnosis confirms what the math above predicts, the next step is already waiting for you at the link below. Every day you defer it costs the same \$12,500.

SEAS Programme Time Horizon — 18 Months vs 24 Months

The SEAS implementation runs in two segments:

18-month programme (Phases 0–4): Builds, deploys, and validates agents through five go/no-go gates. Delivers \$4M–\$6M of cumulative captured EBITDA and an annual run-rate of \$4M–\$5.5M by the end of month 18. This is the SEAS deployment horizon.

24-month programme (adds Phase 5, months 19–24+): Layers bonus systems on top — sales, customer service, compliance, and ecosystem effects. Lifts the cumulative captured EBITDA to \$6M–\$10M and the annual run-rate to \$6M–\$10M by the end of month 24+.

When this Playbook references the headline "\$6M–\$10M" figure, it is the 24-month outcome (Phases 0–5 complete). The 18-month outcome is \$4M–\$6M cumulative. The two figures are sequential milestones, not interchangeable claims.

Start Running FLOAT.

The full SEAS System (309-page playbook, 9 Excel models, 35 agent specs, 20+ legal templates, board decks, software) is one click away.

► **GET THE FULL SEAS SYSTEM — \$2,997** ◀

<https://smartagenticssystems.com/>

90-day 100% refund guarantee — no questions asked. If FLOAT doesn't identify at least 3 actionable EBITDA opportunities, that alone is reason enough.

A note on the \$150M reference company

The math throughout this book is anchored to a \$150 million revenue mid-market company with 12% EBITDA margin — chosen as the median of the SEAS target range (\$50M–\$500M revenue) for ease of calculation. The \$5.3M+ recoverable figure is the diagnostic result for that anchor case. Your firm's recoverable EBITDA scales proportionally with revenue: a \$50M firm

should expect \$1.5M-\$2M; a \$75M firm roughly \$1.7M-\$2.6M; a \$250M firm \$8M-\$10M; a \$300M firm \$10M-\$11M; a \$500M firm \$15M-\$20M.

The diagnostic methodology, the three pockets, and the 14-day pilot path apply identically across the entire range — only the dollar amounts scale.

Chapter 1 — The \$5.3 Million Hiding in Your P&L

Take a \$150 million revenue mid-market company with a 12% EBITDA margin. That is \$18 million of EBITDA on the trailing twelve months. It is the median company in the SEAS target range, which runs from \$50 million to \$500 million in revenue. It is the company most of this book is built around, because it is the company most often left out of the published research on AI value capture. The McKinsey surveys cover Fortune 500 enterprises. The Deloitte reports survey 3,235 senior leaders, predominantly large enterprise. The mid-market \$50M–\$500M segment is footnoted, not focused.

That is a problem, because the mid-market is where the leakage is largest in proportional terms and where the recovery window is shortest. A \$150 million company with the same operational habits as a \$1.5 billion company is leaking the same percentage of revenue, on a much thinner EBITDA base. A 3% revenue leakage on \$150 million is \$4.5 million. On an \$18 million EBITDA base, that is 25% of total profit.

Twenty-five percent of your profit. Leaking out the bottom of the P&L. Every year. While you read transformation articles.

At the upper end of the SEAS range — a \$400M revenue firm with a 12% margin — the same proportional leakage scales to roughly \$32 million of annual waste, with \$14M–\$15M recoverable on a richer base. At the lower end — a \$75M revenue firm — the recoverable figure is \$1.7M–\$2.6M. The \$5.3M figure used throughout this book is the median anchor; the recoverable amount scales with revenue and operational profile across the full \$50M–\$500M range.

That is the prize.

The arithmetic of leakage

Three categories of waste account for almost all of the recoverable EBITDA in a typical mid-market company. They are not glamorous. They are not hidden in some specialised corner of the business. They sit on the P&L in plain sight, and they are large enough that any CFO with a calculator and an honest hour can find them.

Leakage Category	Annual Waste	Recoverable	Capture Method
Supply Chain Inefficiency	\$6.0M	\$2.5M	Predictive demand, autonomous reordering
Legacy Administrative	\$4.0M	\$1.8M	AP/AR automation, RAG

Burden			workflows
Energy and Operations Waste	\$2.0M	\$1.0M	HVAC scheduling, predictive maintenance
Total	\$12.0M	\$5.3M	12-18 months, no capex

The reason this number lands harder than a typical consulting estimate is that each row of the table maps to a process the CFO already touches every month. Supply chain inefficiency is the inventory write-down conversation in the August closing. Legacy administrative burden is the four AP clerks the controller asked to hire last year. Energy and operations waste is the utility bill that nobody on the executive team has the bandwidth to investigate. These are not theoretical leakages. They are the line items the CFO has already noticed and has been told, in three different ways, that there is no time to fix. The diagnostic in this book is what the CFO does on the Tuesday they decide there is time.

\$5.3 million of recoverable EBITDA on \$18 million of baseline EBITDA. That is a 29% lift in profit on no revenue growth, no headcount reduction, and no capital expenditure beyond a \$50,000 to \$150,000 pilot budget.

Read that sentence again. Twenty-nine percent more profit. Without raising prices. Without firing anyone. Without another ERP implementation. For less than the cost of one senior executive's annual bonus.

These are conservative capture rates, drawn from category-level benchmarks published by Accenture, FTI Consulting, BCG, and KPMG between 2024 and 2026. The high-performer figures are higher. Bain's 2025 report on AI value capture documents that AI leaders deliver 10-25% EBITDA gains while laggards see only 1-3%. KPMG's 2025 work on Fortune 1000 labour productivity benchmarks 5% (500 basis points) of EBITDA achievable from AI in labour productivity alone. FTI Consulting's research on AI-centric organisations in the top quartile reports 12 to 14 percentage point EBITDA margin increases versus industry averages.

The SEAS conservative capture case sits below all of these. It assumes a mid-market firm that is not a top-quartile performer, that has not invested in clean data, and that does not have an in-house AI team. It assumes a CFO who is sceptical and a board that wants 90-day proof points before approving Phase 3. It still produces \$5.3 million.

The PE multiple expansion arithmetic

If you are a PE operating partner, the EBITDA number is not the headline. The multiple expansion is the headline.

AI-native portfolio companies — defined as companies with auditable, transferable AI capability documented in the data room and validated by buyer-side diligence — are modeled in the SEAS framework to command 2 to 3 additional turns of EBITDA multiple at exit — consistent with the premium the market has historically paid for demonstrably superior operating capability. An 8x baseline becomes 10x to 11x. On a \$150 million revenue company with \$18 million of EBITDA growing to roughly \$24 million through SEAS deployment (the \$18 million base plus the recovered EBITDA, rounded for this illustration), the enterprise value bridge looks like this:

Scenario	Exit EBITDA	Multiple	Enterprise Value	Value vs Base
Base (No AI)	\$18.0M	8x	\$144M	—
AI-EBITDA Only	\$24.0M	8x	\$192M	+\$48M
AI-Native (EBITDA + Premium)	\$24.0M	10x	\$240M	+\$96M
High-Performer (Top Quartile)	\$24.0M	11x	\$264M	+\$120M

\$120 million of additional enterprise value creation. On a single portfolio company. From a diagnostic that takes 90 minutes to run and a system that costs \$2,997 to buy. If you are running a PE fund with five portfolio companies in the SEAS target range, multiply that number by five. If you are running a fund with ten, multiply by ten. If the math looks implausible at that point, open the MOIC Calculator included in the SEAS package and plug in your own numbers. That is the entire point of the calculator — it is not a marketing claim, it is a live financial model that you can audit line by line.

The EBITDA improvement alone delivers \$48 million of value creation. The multiple expansion on top of that delivers another \$48 million in the base AI-Native case, and another \$72 million in the high-performer case. The multiple expansion is the same magnitude as the EBITDA improvement, and in the high-performer case it is larger. These are simplified, directional enterprise-value figures — EBITDA times multiple, with no debt, organic growth, or paydown modelled. They include both the recovered EBITDA valued at the exit multiple and the multiple re-rating on the existing \$18 million earnings base. SEAS's leakage table reports a narrower figure — a \$53 million

enterprise-value uplift — because it values only the recovered EBITDA at the exit multiple and excludes the re-rating on the base. The two answer different questions; neither contradicts the other.

This is the arithmetic that PE operating partners need to internalise. SEAS is not a cost-reduction programme. It is a multiple-expansion programme that uses cost reduction as the proof mechanism. The board does not buy AI because it saves money on invoices. The board buys AI because it adds two turns to the exit multiple, and the saved money on invoices is the audit trail that justifies it.

Why this opportunity is open right now

Three things have converged to make this the moment, and they will not stay converged forever.

First, agentic AI has crossed the production threshold. Deloitte's 2026 Tech Trends report documents that 11% of organisations have deployed agentic AI in production, up from a rounding error in 2024. The technology works. It is not experimental. The vendors have shipped. Microsoft Copilot, UiPath, Automation Anywhere, Moveworks, Celonis — each has documented case studies with named customers, named outcomes, and named ROI multiples. The risk of being early is now lower than the risk of being late.

Second, the implementation cost has collapsed. A pilot deployment that would have cost \$500,000 in 2022 now costs \$50,000 to \$150,000. The vendors have rationalised their pricing because they want production references, and the open-source ecosystem has commoditised the underlying components. Most mid-market companies can fund a pilot out of operating budget, with no board approval and no capital request.

Third, the gap between leaders and laggards is widening. McKinsey's 2026 research on PE GP AI deployment reports that only 6% of GPs see high AI impact today, but 70% expect significant impact in 3–5 years. That is the arbitrage window. The PE firms that move now will sell their AI-native PortCos to the PE firms that wake up in 2028 and realise they have been left behind. The buyer pool for AI-native exits is about to expand by an order of magnitude. The seller pool is still small.

The arbitrage window is roughly 24 to 36 months. After that, AI-native operations become the default expectation in mid-market PE diligence, and the premium evaporates. By 2028, you will not get extra turns for having AI. You will get a discount for not having it.

Every month of delay on a \$150 million revenue company with 3% addressable EBITDA leakage costs roughly \$375,000 in unrealised margin. That is the run-rate cost of waiting. Six months of waiting is \$2.25 million. Twelve months of waiting is \$4.5 million. None of it is recoverable. It is gone the moment the calendar turns.

That is the case for running FLOAT this week. Not next quarter. Not after the budget cycle. This week.

The Math Is Undeniable. The Next Step Is One Click.

SEAS gives you the full MOIC Calculator, the 18-month roadmap, the 20+ legal templates, and the board deck. Download it and audit every number yourself.

▶ GET THE FULL SEAS SYSTEM — \$2,997 ◀

<https://smartagenticssystems.com/>

90-day 100% refund guarantee — no questions asked. If FLOAT doesn't identify at least 3 actionable EBITDA opportunities, that alone is reason enough.

Chapter 2 — Why Three Out of Four AI Initiatives Fail

RAND's 2024 study on enterprise AI project outcomes is the most uncomfortable piece of research published on this topic. It found that 65% to 80% of enterprise AI initiatives fail to deliver the expected ROI. Not 20%. Not 30%. Most of them fail.

This is the number that should sit at the front of every board memo on AI investment. It does not mean that AI does not work. The technology works. It means that the way most organisations deploy AI does not work, and the failure modes are predictable enough that you can point at them in advance.

There are three of them. They account for the overwhelming majority of failures. Each one has a single root cause. Each one has a single antidote. The antidotes are not technical. They are organisational. And every one of them is built into the SEAS system as a structural safeguard — which is the reason the PE operating partners using SEAS are not part of the 65–80% failure statistic.

Failure Mode 1 — Wrong Vendor, Wrong Problem

Frequency: roughly 30% of AI failures.

This is the failure mode that begins with a vendor demo. A salesperson from a well-funded AI startup walks into the office, shows a slick demonstration of a chatbot answering questions about a HR policy, and the CIO walks out of the meeting convinced that the company needs to deploy AI in HR. Six months and \$400,000 later, the company has a working chatbot that handles 200 HR queries per month, no measurable EBITDA impact, and a CFO asking what just happened.

The root cause is that the deployment started with the technology, not the business problem. The vendor knew its product. The vendor did not know which P&L line was leaking. So the deployment optimised the line that was easiest for the vendor to address, not the line that was bleeding the most cash.

The antidote is to start with the diagnostic. Identify the EBITDA pockets first. Rank them by size and speed-to-cash. Then go to vendors with a defined problem and ask which of them can solve it. Vendors compete to win a defined problem. Vendors do not compete to define a problem on your behalf — they will define the problem in the shape of their product, every time.

FLOAT exists to prevent this failure mode. The 90 minutes you spend on the diagnostic is the cheapest insurance policy in enterprise software. It costs you a coffee. It saves you \$400,000.

Failure Mode 2 — Scope Creep

Frequency: roughly 37% of AI failures. This is the largest category.

Scope creep starts the same way every time. A pilot succeeds. The team is excited. Leadership wants to capitalise on the momentum. Three more workflows are added to the pilot before the first one has stabilised. A fourth is added before the second has shown ROI. By month nine, the original pilot is one of seven concurrent deployments, none of them are at production accuracy, and the steering committee is meeting every two weeks to triage exceptions instead of measure outcomes.

The root cause is the absence of phase gates. A pilot that has not been validated is not a pilot. It is an experiment. Experiments do not scale. Validated pilots scale. The discipline is to insist on validation — measured against pre-defined criteria, not against vibes — before any scaling decision is made.

The antidote is the go/no-go gate. SEAS uses five of them, one between each phase. The most important is Gate 1, at Week 8, which gates the transition from pilot to production. The criteria are not negotiable. Agent accuracy must exceed 90% on a manually validated sample of at least 100 transactions. Cycle time must be reduced by at least 25% versus baseline. Cost per transaction must be reduced by at least 15%. Exception rate must be below 8%. User satisfaction must be above 3.5 out of 5. Stable operation for at least 4 weeks at 95% uptime. Data quality issues unresolved must be below 2%. Zero unplanned downtime in the prior two weeks. Team trained and confident. Budget variance under 10%.

Ten criteria. Pass threshold is at least 7 of 10. If fewer than 7 are met, the pilot pauses. It does not scale. It does not get a second chance to scale next month. It pauses. Two weeks of remediation. Re-evaluation. Then a decision.

This sounds harsh. It is the reason SEAS-implemented pilots scale and most other pilots do not. The discipline is not about being conservative. The discipline is about being honest. A pilot that cannot meet 7 of 10 criteria after 8 weeks is a pilot that was deployed against the wrong workflow, the wrong data, or the wrong vendor. Scaling that pilot does not fix the problem. It compounds the problem.

Failure Mode 3 — No Accountability

Frequency: roughly 33% of AI failures.

This one is the simplest to describe and the hardest to fix. The pilot has no named owner. The steering committee meets once a quarter. KPIs are measured monthly, then quarterly, then not at all. Nobody on the executive committee can answer the question "how is the AI pilot doing" without consulting three different people, none of whom actually own the outcome.

The root cause is that AI sits in the middle of the org chart. It is not IT. It is not Finance. It is not Operations. It touches all three, and in most companies, that means it is owned by none of them.

The antidote is to put the CFO on the hook. Not the CIO. Not the COO. The CFO. The reason is structural. The CFO is the only executive whose performance is measured directly in EBITDA, which is the metric AI is supposed to move. If the CFO sponsors the deployment, the CFO has skin in the game. The metric the CFO reports to the board is the metric the AI deployment is judged against.

SEAS adds a second piece of accountability infrastructure: the CEO Alignment Charter. This is a board-approved document that links 20% to 30% of the CEO's Year 1 bonus to AI-driven EBITDA milestones. It is a legal template included in the SEAS package. It exists because PE operating partners learned, expensively, that voluntary CEO commitment to AI deployment evaporates the moment the next quarter's revenue target slips. Compulsory CEO commitment, written into the bonus structure, does not.

Three failure modes. Three antidotes. The diagnostic. The phase gate. The CFO sponsorship. None of them are technical. All of them are organisational. AI failure is almost never an AI problem. It is a project management problem dressed up in a hoodie. SEAS is built to solve the project management problem — which is why it works where consulting engagements and vendor pilots fail.

Chapter 2.5 — The 2026 Mid-Market Operating Reality

The published research on AI value capture is dominated by Fortune 500 case studies. There is a structural reason for this, and the reason is not flattering to the research community: Fortune 500 companies have research budgets, communications teams, and PR cycles that turn pilot deployments into published case studies. Mid-market companies do not. A \$200 million revenue distributor that captures \$7 million of EBITDA from agentic AI does not write a McKinsey article about it. The CFO files a board update, the PE sponsor counts the EBITDA, and the company keeps operating. The case study never enters the public record.

This creates a strange asymmetry in 2026. The companies most able to deploy agentic AI quickly — the mid-market \$50M to \$500M segment, where decision cycles are short and operational sponsorship is concentrated in three or four executives — are the segment with the least published guidance. The Fortune 500 case studies are not portable. A 50,000-employee enterprise-wide deployment of an AI workflow tool tells a 600-employee mid-market manufacturer almost nothing about how to deploy the same tool in their own operation. The complexity, the governance overhead, and the integration challenges scale non-linearly.

What FLOAT addresses, at its core, is the absence of mid-market-specific operational guidance. The three pockets in this book — SG&A efficiency, pricing drift, and supply chain — are the three pockets that show up consistently in the \$50M to \$500M segment, in the proportions documented in Chapter 1. They show up because of how mid-market companies are structured, not because of any particular industry or geography. A regional medical-device distributor in the Midwest, a specialty chemical manufacturer in the United Kingdom, a B2B logistics platform in Singapore, and a packaged-food brand in southern Europe will all show variants of the same three pockets. The pockets are not a function of the business model. They are a function of the operating reality of running a \$50M to \$500M company without a fifty-person McKinsey team on retainer.

Why the leakage persists

The leakage persists for three structural reasons.

First, mid-market CFOs are full-time on cash, debt covenants, and quarterly close. The CFO of a \$200 million company has a controller, an FP&A manager, an accounts payable lead, and an accounts receivable lead. That is roughly the entire finance function. Nobody on that team is full-time on margin recovery. Margin recovery is what the consultants are supposed to do, except the consultants are too expensive to engage on a continuing basis.

Second, the operational data needed to run a structured diagnostic exists, but is fragmented across three or four systems that do not speak to each other in real time. The ERP has the P&L. The procurement system has the vendor concentration. The CRM has the contract terms. The HR system has the headcount-to-revenue ratios. Pulling all four into a single view takes either an analyst with a week of bandwidth or a structured methodology like FLOAT that uses the existing reports rather than requiring new ones.

Third, and most important, the operating cadence of a mid-market company does not have a slot for "run the structured EBITDA leakage diagnostic." There is no Q3 budget cycle line item for it. There is no five-year strategic plan section for it. There is no executive committee agenda item for it. The work happens because somebody on the executive team — typically the CFO or a PE operating partner stepping in — decides to schedule 90 minutes on a specific Tuesday and run the diagnostic. That decision is the entire intervention.

What changed in 2024 to 2026

For the first twenty years of corporate AI initiatives, the right answer for a mid-market CFO was usually to wait. The technology was not mature. The vendors were unstable. The implementation costs were prohibitive. The case studies were inflated. A patient CFO who waited could let the early-adopter pain land on Fortune 500 budgets while learning from the public failures.

That window closed in 2024. The current generation of agentic AI tools — predictive AP automation, contract value capture agents, demand forecasting with autonomous reordering — has shipped, with mid-market deployments, transparent pricing, and standard ERP integrations. Implementation costs that ran \$2M to \$5M in 2018 now run \$50K to \$150K for a comparable pilot. Deployment timelines that ran 18 to 24 months now run 8 to 12 weeks to first value.

The CFOs who are still waiting in 2026 are no longer waiting for technology maturity. They are waiting for permission. Permission from a board that has not seen the case study yet. Permission from a CEO who is focused on top-line growth. Permission from a PE sponsor who has not raised the question. Permission from an industry peer group that is also waiting.

The waiting is the cost. The cost is \$375,000 a month for a \$150 million revenue company, and it scales linearly with revenue across the SEAS target range.

FLOAT exists for the CFO or PE operating partner who has decided to stop waiting. The 90-minute diagnostic gives you the dollar number, named in your firm's actual P&L, that you can take to the next board meeting. If the dollar number is large enough — and on a \$150 million company it is almost always between \$4 million and \$7 million — the conversation about whether to act is not a conversation any longer. The conversation becomes which pocket to attack first.

Chapter 3 — FLOAT: The Framework

FLOAT stands for Firm Leakage Opportunity And Triage. The acronym is engineered to do real work. Each letter names a phase. Each phase has a deliverable. Each deliverable has a go/no-go gate.

Phase	What Happens	Deliverable	Time
F — Find	Identify EBITDA leakage pockets via data forensics	Three pockets named with red flags identified	30 min
L — (Leakage)	(Embedded across all phases — what the diagnostic identifies)	—	—
O — Opportunity	Quantify waste, map root causes, size the dollar opportunity	Three pockets sized in dollars with capture %	30 min
A — And	(Connector — diagnostic produces both opportunities and the triage to act on them)	—	—
T — Triage	Rank by speed-to-cash; pick one pilot workflow	Named pilot, named owner, named demo, named vendor	30 min

The active phases are Find, Opportunity, and Triage. Leakage and And are connectors that complete the acronym and signal what the diagnostic is operating on. Three active phases of 30 minutes each gives you the 90-minute diagnostic. Two short coffee breaks between phases is reasonable. Less than two is heroic. More than two and you are not doing FLOAT, you are doing a workshop.

What you need before you start

Three documents. That is the entire prerequisite list. If you cannot pull them in fifteen minutes, your data infrastructure is the actual problem and FLOAT will tell you that within the first 30 minutes, which is also useful information.

- The trailing twelve-month profit and loss statement. Standard format. Revenue, COGS, gross profit, SG&A by line item, operating income, EBITDA. Most CFOs have this open in a tab right now.
- The SG&A breakdown by department. Finance and Accounting, Customer Operations, IT and Helpdesk, Supply Chain Operations, HR, Legal, Sales and Marketing G&A. Each as a dollar figure and a percentage of revenue.

- The vendor concentration report. Top 20 vendors by spend. Total annual spend per vendor. Date of last competitive bid. Auto-renewal status.

That is it. No new data collection. No special software. No consultant. The point of the diagnostic is that the data already exists in every mid-market company that produces a quarterly board pack. FLOAT just looks at the existing data through a specific lens — the lens of recoverable EBITDA — that most board packs do not use.

The 90-minute structure

Set aside 90 minutes. Block the calendar. Close the laptop on email. Sit in a room with the three documents and a single piece of paper. Across the top of the paper, write three column headings: Pocket 1 — SG&A Efficiency Gap. Pocket 2 — Pricing and Contract Drift. Pocket 3 — Supply Chain and Vendor Concentration. Down the left side, write four row headings: Red flag check. Dollar quantification. Root cause hypothesis. Next-action workflow.

That four-by-three grid is the entire deliverable of the FLOAT diagnostic. Twelve cells. Filled in with specific numbers, not adjectives. At the end of 90 minutes, you have a single page that names three EBITDA pockets, sizes each one in dollars, identifies the root cause of each, and points at a specific workflow that an agentic AI deployment could attack.

From that one page, you pick the pocket with the largest dollar opportunity that has the lowest implementation complexity. That is your pilot. You then have an action item: schedule a vendor demo by Friday.

Chapters 4, 5, and 6 walk through each of the three pockets in turn. Each chapter has the red-flag thresholds, the dollar quantification formulas, and the root cause questions. Read them once. Then run the diagnostic with the documents in front of you.

Chapter 4 — Pocket 1: The SG&A Efficiency Gap

The first 30 minutes of the diagnostic. This is the largest pocket in most mid-market companies, and it is the one that usually surprises the CFO the most, because it sits in plain sight on the SG&A schedule that the CFO signs every month.

The reason it is invisible is that every individual line item looks reasonable in isolation. Finance and Accounting at 7.7% of revenue is not glaringly out of line for a \$150 million company. Customer Operations at 5.7% looks fine. IT and Helpdesk at 2.4% is conservative. Add them up and the total SG&A is 24% of revenue, which is in the normal range for a mid-market manufacturer or distributor.

The problem is that "in the normal range" is benchmarked against the median. The median includes the laggards. The best-in-class number for Finance and Accounting is not 7.7% of revenue. It is closer to 1.5%. That is more than a six-point gap on a single line item, which is roughly \$9 million on a \$150 million company. The recoverable portion of that gap — the portion that can be captured with current agentic AI tools, conservatively, in 12 months — is around \$2.2 million.

That is one line item. Pocket 1 has five red-flag checks.

The five red-flag checks (30 minutes)

Pull the SG&A breakdown. Run the following five checks in order. Each one takes about five minutes. The sixth five minutes is for adding up the dollar figures and writing them on the diagnostic worksheet.

Check	Red Flag	EBITDA Impact
SG&A as % of revenue versus peer median	Exceeds peer median by ≥ 2 percentage points	2 pp gap on \$150M = \$3M opportunity
Finance and Accounting as % of revenue	>3% of revenue (best-in-class is $\approx 1.5\%$)	Target: <1.5%, gap recoverable in 12 months
Cost per invoice processed	>\$3 to \$4 per invoice	50K invoices $\times$ \$2.50-\$4.50 savings = \$125K-\$225K
AR DSO versus industry median	>45 days with no automation in place	10-day DSO cut = \$4.1M one-time WC release (not annual EBITDA)

AP/AR FTEs per \$100M of revenue	>8 FTEs per \$100M (benchmark is 4-6)	FTE reduction at \$80K loaded cost each
----------------------------------	---------------------------------------	---

Take the dollar figures from the third column. Add them. Write the sum on the diagnostic worksheet next to "Pocket 1 — Annual EBITDA Opportunity." For a typical \$150 million revenue company that has not addressed any of these five issues, the number lands between \$2 million and \$5 million. That is the size of Pocket 1.

The root cause and the next-action workflow

The root cause of the SG&A efficiency gap is almost always the same in mid-market companies, and it is structural rather than cultural. The finance team has grown linearly with revenue because every new million in revenue brings new invoices, new collections, new vendor reconciliations, new month-end close activities. Every CFO has had the same conversation with the finance director three times: we cannot close the books faster without hiring more people, and we cannot hire more people without budget approval. The hiring happens. The cost goes up. The work continues to be done by humans clicking through ERP screens.

Agentic AI breaks the linearity. A predictive AP automation deployment processes 80% of invoices end-to-end without human touch. The 50,000 invoices that previously consumed four AP clerks now consume one AP clerk plus the agent. The cost per invoice drops from \$4.50 to \$1.80. SG&A as a percentage of revenue declines mechanically, year over year, without anyone being fired.

If Pocket 1 turns up the largest dollar figure on your diagnostic worksheet, the pilot workflow is almost certainly AP automation — the workflow with the cleanest baseline data, the fastest deployment timeline (8 to 10 weeks), the most mature vendor ecosystem, and the most defensible ROI calculation. The full SEAS package includes the vendor matrix scoring ten agentic AI platforms across six weighted criteria, plus the complete Agent Prompt Library for validating the opportunity in 15 minutes with your own data before committing budget.

Illustration — Pocket 1 in a Hypothetical \$200M Distributor

To make the SG&A diagnostic concrete, walk through the math for a hypothetical \$200 million revenue B2B distributor with a 10% EBITDA margin — a representative case in the SEAS target range. The numbers below are illustrative, not drawn from a specific client engagement, but the proportions track what a CFO with the standard reports in front of them will typically see.

The Finance and Accounting line on the SG&A schedule sits at \$9.4 million, which is 4.7% of revenue. Best-in-class for a distributor of this size is closer to 1.8%. The gap is

2.9 percentage points, or roughly \$5.8 million of structural overhead. The recoverable portion — the portion that can be captured in 12 months with current AP automation tools — is conservatively 25 to 30%, which produces a Pocket 1 contribution of \$1.5 million to \$1.7 million from this single check.

The cost-per-invoice check produces a second contribution. The distributor processes roughly 65,000 invoices per year at an internal cost of \$4.20 per invoice. The agent-assisted target is \$1.80. The savings of \$2.40 per invoice across 65,000 invoices is \$156,000 of annualised gross EBITDA. Modest, but it stacks on top of the F&A gap.

DSO is 54 days against an industry median of 41. The 13-day cut releases approximately \$7.1 million of working capital — a one-time release that lands on the cash-flow statement, not annual EBITDA. The CFO records this on the diagnostic worksheet as a separate line, not added to the recurring savings number, because conflating the two creates the kind of confusion that makes IC memos lose credibility.

Total Pocket 1 for this hypothetical firm: roughly \$1.7 million to \$1.9 million of annual recurring EBITDA, plus the \$7.1 million one-time WC release. Captured in 12 months. Pilot budget under \$150,000. The candidate workflow is predictive AP automation, deployable in 8 to 10 weeks, with payback inside 6 months at full production volume.

That is the answer the diagnostic produces in 30 minutes. The CFO writes the numbers in the Pocket 1 cells of the worksheet and moves on to Pocket 2.

Chapter 5 — Pocket 2: Pricing and Contract Drift

The second 30 minutes. This pocket is smaller in absolute dollars than Pocket 1 in most mid-market companies, but it has the highest margin flow-through and the fastest payback. Pricing recovery is the only EBITDA lever that flows 100% to operating profit on day one with no implementation lag. The day you reprice a contract from a 15% discount to a 12% discount is the day the additional margin starts arriving.

The reason this pocket exists is universal and embarrassing. Sales teams negotiate discounts to close deals. The discounts get baked into the contract. The contract auto-renews. Nobody re-prices it. Costs go up. Inflation arrives. The discount that was 8% in 2020 is now effectively 18% in 2026 because the input costs have risen and the price has not. Nobody notices because nobody is looking at the right report.

AI does not fix this problem in some mystical way. AI fixes this problem by reading every contract in the customer base, extracting the pricing terms, comparing them against current cost, and producing a list of contracts that are leaking margin. A human still has to make the renegotiation decisions. The agent just makes sure the human knows where to look.

The four red-flag checks (30 minutes)

Check	Red Flag	EBITDA Impact
% of contracts not repriced in 18-24 months	$\geq 10\%$ of revenue on legacy pricing	0.5-3% margin recovery = \$750K-\$4.5M
SKU margin variance across product lines	$\geq 5\%$ spread without active management	Bottom 20% SKUs are repricing targets
Average discount depth versus list price	$> 15\%$ average discount from list price	Discount leakage without enforcement
Auto-renewal contracts without escalation	Renewals lacking CPI or cost-pass-through	Repricing trigger identified

Pull the customer master. Filter for contracts with a last-modified date older than 18 months. Sum the annual contract value of that filtered list. If it exceeds 10% of total revenue, you have the pricing drift problem. Multiply that ACV by an assumed 3% recoverable price increase and you have the dollar size of the opportunity. For a \$150 million company with \$30 million of legacy-priced revenue, that is \$900,000

minimum. The ceiling, in companies with severe pricing neglect, is closer to \$3 million.

The cumulative cost of leaving pricing alone for three years on a \$150 million company with normal cost inflation is roughly \$4 million in unrecovered margin. That number is not hypothetical. That number is the math of a P&L running on autopilot while nobody on the executive team is full-time on margin recovery. The CFO does not own pricing. The CRO owns pricing, and the CRO is measured on revenue retention, not margin recovery. So the margin leaks.

Find the \$4 Million Your Contracts Are Leaking.

The full SEAS package includes the Contract Value Capture agent specification plus the Templates Library with ready-to-use renegotiation frameworks.

► **GET THE FULL SEAS SYSTEM — \$2,997** ◀

<https://smartagenticsystems.com/>

90-day 100% refund guarantee — no questions asked. If FLOAT doesn't identify at least 3 actionable EBITDA opportunities, that alone is reason enough.

Illustration — Pocket 2 in a Hypothetical \$120M B2B Software Firm

Pricing drift is most visible in B2B software, professional services, and contractual subscription businesses, where revenue is concentrated in a small number of large contracts that auto-renew without active price re-examination. Walk through the math for a hypothetical \$120 million revenue B2B software company with a 22% EBITDA margin. Again, the numbers are illustrative, not drawn from a specific engagement.

Pull the customer master. Filter for contracts with a last-modified date older than 24 months. The filtered list captures \$34 million of annual contract value, which is 28% of total revenue. That number alone exceeds the 10%-of-revenue threshold for Pocket 2 by a wide margin. The pricing drift problem is real and it is sitting in the customer master, untouched, for the simple reason that nobody on the executive team is full-time on margin recovery.

Apply a conservative 3% recoverable price increase on the legacy-priced subset. That produces \$1.0 million of annual EBITDA recovery at 100% margin flow-through. In a

software business, where every additional dollar of price runs straight to operating profit, this is a uniquely clean pocket — there is no implementation cost, no vendor licensing fee, no integration risk. The work is renegotiation, executed contract by contract, supported by an agent that reads each contract, extracts the pricing terms, compares them against current cost trajectories, and produces a list of contracts that are leaking margin.

The ceiling is higher. A 6% effective price increase on the legacy subset, achievable in software businesses where prices have not been adjusted for three or more years, produces \$2.0 million. Combined with a CPI escalation clause for renewals lacking one, the total Pocket 2 contribution lands at \$1.5 million to \$2.2 million in this hypothetical case.

The sales team will resist. The CRO is measured on revenue retention, not margin recovery. The repricing conversation is uncomfortable. None of those facts changes the underlying arithmetic. The contract that was priced for 2022 cost levels is, in 2026, leaking margin every quarter. Either it is renegotiated or it continues to leak. There is no third option.

The candidate pilot for Pocket 2 is a Contract Value Capture agent — payback period 6 to 12 months, implementation cost typically below \$80,000, and 100% margin flow-through from day one of any successfully renegotiated contract. The diagnostic worksheet records the dollar opportunity, the candidate workflow, and the named pilot owner. Move to Pocket 3.

Chapter 6 — Pocket 3: Supply Chain and Vendor Concentration

The third 30 minutes. This pocket is the largest in absolute dollars in industrial, manufacturing, and distribution companies, and it is the one with the longest payback period. Pocket 1 pays back in three to six months. Pocket 2 pays back in six to twelve. Pocket 3 pays back in nine to eighteen, and the deployment is more complex because it touches multiple systems and multiple vendor relationships.

None of that is a reason to skip it. Pocket 3 is also the pocket with the highest absolute dollar value, the most defensible at exit, and the most attractive to PE buyers in the diligence process. A demand forecasting agent that has reduced inventory by 15% with no service-level degradation is a story that sells exit multiples. An invoice processing agent is a story that sells operating efficiency. Both are valuable. The first one is more valuable at exit.

The four red-flag checks (30 minutes)

Check	Red Flag	EBITDA Impact
Top 3 suppliers as % of total COGS	>60% supplier concentration	1-3% margin = \$1.5M-\$4.5M opportunity
Last competitive bid for top vendors	No competitive bid in 2+ years	0.5-1% COGS savings recoverable
Inventory turns versus peer median	>15% below peer median	15% inventory reduction = ~\$3M working capital release
Energy cost variance year over year	>10% YoY without volume explanation	8-15% energy cost reduction achievable

If Pocket 3 turns up the largest dollar figure on your diagnostic worksheet, the candidate pilot workflow is almost always demand forecasting with autonomous reordering — the workflow with the highest absolute dollar impact, the strongest exit narrative, and the most defensible ROI math against a finance audience. The deployment is more complex than AP automation (Pocket 1's typical pilot) and the payback is longer (9-18 months versus 3-6), but the ceiling is higher and the diligence story at exit is stronger. If your firm is heading toward a sale or refinancing in the next 18-24 months, this is the pocket whose workflow you want documented in the data room.

Pull the vendor concentration report. Identify the top three suppliers by annual spend. Sum their spend as a percentage of total COGS. If the answer is over 60%, you have the supplier concentration risk. Pull the inventory turns metric. Compare to peer median for your industry. If you are 15% below peer median, you are carrying excess inventory, which means working capital tied up in stock that could be released. For a \$150 million revenue company with a 13% inventory-to-revenue ratio, a 15% inventory reduction is a \$2.9 million one-time working capital release plus the carrying cost savings.

Pull the energy bills. Compare year over year. If energy costs have risen more than 10% without a corresponding volume increase, you are paying for inefficiency. Predictive HVAC scheduling and demand-response optimisation typically deliver 8 to 15% energy cost reduction in industrial and distribution facilities, on payback periods of 12 to 18 months.

Supply chain leakage is the only one of the three pockets where the root cause is genuinely complex. Pocket 1 is structural understaffing automation. Pocket 2 is misaligned incentives. Pocket 3 is the cumulative effect of dozens of operational decisions made under uncertainty over a period of years. The supply chain that was carrying eight weeks of safety stock to cover a 2021 supplier disruption is still carrying eight weeks of safety stock in 2026, even though the disruption ended four years ago. Agentic AI is the only technology with the analytical bandwidth to look at all the stock layers simultaneously and ask whether each one is still justified.

Illustration — Pocket 3 in a Hypothetical \$250M Industrial Manufacturer

Supply chain leakage is largest in absolute dollars in industrial, manufacturing, and distribution firms, and the diagnostic produces the most consequential exit-narrative findings here. Walk through the math for a hypothetical \$250 million revenue industrial manufacturer with a 14% EBITDA margin. Recoverable EBITDA at this revenue scale is in the \$8 million to \$10 million range across all three pockets, with Pocket 3 typically the largest contributor.

Pull the vendor concentration report. The top three suppliers account for 68% of total COGS. That exceeds the 60% threshold and signals concentrated negotiating exposure. None of the top three has been competitively bid in the last 30 months. Apply a conservative 1.5% COGS savings target — a defensible recovery from a competitive rebid for vendors that have not faced market pricing pressure in two or more years. On a COGS base of roughly \$175 million, that is \$2.6 million of annual recurring EBITDA.

Pull the inventory turns metric. The firm runs 4.1 turns against a peer median of 5.2 — roughly 21% below median. Inventory of \$42 million sits on the balance sheet, against

a peer-aligned target of \$33 million. A 15% inventory reduction releases \$6.3 million of working capital plus carrying-cost savings of approximately \$300,000 annually. The working-capital release is one-time. The carrying-cost savings recur.

Energy costs have risen 14% year over year against volume growth of only 4%. The 10-percentage-point gap signals operational inefficiency in scheduling, HVAC, or production planning. Predictive scheduling and demand-response optimisation typically deliver 8 to 12% energy cost reduction in industrial facilities of this size. On a \$4.8 million annual energy spend, that is \$385,000 to \$575,000 of annual EBITDA recovery.

Total Pocket 3 for this hypothetical manufacturer: \$3.3 million to \$3.5 million of annual recurring EBITDA, plus the \$6.3 million one-time working capital release. Pocket 3 alone, at this revenue scale, is roughly the size of the entire \$5.3 million leakage figure used in the \$150M reference case throughout this book — which is exactly what the linear scaling in Chapter 1 predicts.

The candidate workflow is demand forecasting with autonomous reordering, deployed in parallel with a vendor rebid programme. Payback is 9 to 18 months, longer than Pocket 1 or Pocket 2, but the absolute dollar value is largest, and the data-room narrative for an exit in 18 to 24 months is the strongest of the three pockets. If this hypothetical firm is heading toward a sale or refinancing in that window, Pocket 3 is the pilot.

The diagnostic worksheet now has all three pockets sized. Move to Chapter 7 to pick one.

Chapter 7 — Sizing the Prize: Your ROI Calculator

Ninety minutes after you started the diagnostic, you have a worksheet with three pockets, each sized in dollars, each with a named root cause and a named candidate workflow. You also have a decision to make. You can pursue all three. You should not, at least not yet. The discipline of starting with one pilot, validating it, and only then scaling to the second, is the discipline that separates the 20–35% of AI deployments that succeed from the 65–80% that fail. Pick one. The ROI calculator in this chapter will tell you which one.

This is the same calculator the SEAS Playbook uses to underwrite its pilot recommendations — the live version in the full SEAS package has 1,711 formulas across 9 Excel workbooks and lets you audit every calculation line by line. The version in this chapter is the 13-line paper version. Run it three times. Compare. Pick.

The 13-line ROI calculator

Line	Item	Formula	Example (\$150M Co., AP Pilot)
A	Current annual revenue	—	\$150,000,000
B	Current EBITDA	—	\$18,000,000
C	Current EBITDA margin	$B \div A$	12.0%
D	Pilot workflow baseline cost per transaction	From process map	\$4.50 per invoice
E	Target cost per transaction (post-agent)	Vendor estimate	\$1.80 per invoice
F	Savings per transaction	$D - E$	\$2.70
G	Annual transaction volume	From process data	50,000 invoices
H	Annualised gross EBITDA from pilot	$F \times G$	\$135,000
I	Implementation cost (one-time)	Vendor + internal	\$150,000
J	Annual licensing cost	Vendor quote	\$80,000
K	Year 1 net EBITDA impact	$H - I - J$	\$-95,000 (Year 1 net)
L	Year 2 net EBITDA impact	$H - J$	\$55,000

	(no impl. cost)				
M	Payback volume)	period	(full	$(I + J) \div (H \div 12)$	20.4 months

That example, deliberately, is a small AP pilot — one workflow, one department, 50,000 invoices a year. The economics are positive but unimpressive at this scope, and they are meant to be. The pilot is the proof point, not the prize. A negative Year 1 net of \$95,000 against a \$135,000 annualised gross is exactly what an honest pilot looks like before the implementation cost is amortised across full production volume. Year 2 turns solidly positive without any further investment.

The number that matters is what happens when the validated pilot scales. The same calculator applied to the full enterprise scope — 200,000 invoices a year across the company plus expense reports plus vendor reconciliation — produces approximately \$1.2 million of annualised EBITDA on \$300,000 of incremental implementation cost. Year 1 ROI on incremental implementation cost is 4x. Payback is four months. That is the figure on the IC memo. The pilot's \$95,000 loss is the entrance fee that earns the right to deploy at scale, and the \$1.2 million is what the entrance fee buys.

This is also why the SEAS Playbook reports cumulative figures of \$4M-\$6M by month 18 and \$6M-\$10M by month 24 — those numbers are not optimistic projections, they are the additive result of three or four scaled deployments running in parallel after their respective pilots have been validated. The 13-line calculator in this chapter is the economics of one of those deployments in isolation. The full programme arithmetic is what you find in the SEAS Playbook when you are ready for it.

Walking Through a Filled Worksheet

Chapters 4 to 7 give you the methodology, the red-flag thresholds, and the ROI calculator. This chapter shows you what a completed worksheet actually looks like for the canonical \$150 million revenue mid-market reference company that anchors this book. Read it once. Then run your own.

The reference firm is a mid-market industrial distributor with \$150 million in revenue, \$18 million in EBITDA, and a 12% EBITDA margin. The CFO has 90 minutes, a printed worksheet, and the three required documents — the trailing twelve-month P&L, the SG&A breakdown by department, and the vendor concentration report. The worksheet output below is what the same CFO produces by lunch.

Phase 1 — Find: Pocket 1 (SG&A Efficiency Gap)

The SG&A breakdown shows Finance and Accounting at \$10.2 million, which is 6.8% of revenue. Best-in-class is 1.5%. The gap is 5.3 percentage points. The recoverable portion in 12 months is conservatively 25%, producing \$2.0 million of annual EBITDA. The cost-per-invoice check produces an additional \$135,000 from automating the firm's 50,000 annual invoice flow. AR DSO at 52 days versus an industry median of 42 days releases \$4.1 million of one-time working capital — recorded separately from recurring EBITDA.

Pocket 1 dollar quantification cell: \$2.1 million annual recurring + \$4.1 million one-time WC release. Root-cause hypothesis: headcount grew linearly with revenue; AP and AR processes were never automated. Candidate workflow: predictive AP automation. Implementation timeline: 8 to 10 weeks to Gate 1. Pilot budget: \$120,000 to \$180,000. Payback: 4 to 6 months at full enterprise volume.

Phase 2 — Opportunity: Pocket 2 (Pricing and Contract Drift)

The customer master shows \$27 million in revenue running on contracts last modified more than 18 months ago — 18% of total revenue. Apply a 3% recoverable price increase: \$810,000 of annual EBITDA at 100% margin flow-through. Add the SKU margin variance check: the bottom-quintile SKUs by margin are concentrated in three product lines that have not had a price review in two years; targeted repricing produces an additional \$400,000 of annual recovery.

Pocket 2 dollar quantification cell: \$1.2 million annual recurring. Root-cause hypothesis: discounts negotiated to close deals were baked into auto-renewing contracts; the CRO is measured on revenue retention rather than margin recovery; nobody is full-time on margin. Candidate workflow: a Contract Value Capture agent

that reads every contract in the customer base, extracts pricing terms, compares against current cost trajectories, and produces a renegotiation queue. Implementation timeline: 6 to 8 weeks. Pilot budget: \$40,000 to \$80,000. Payback: 3 to 6 months.

Phase 2 — Opportunity: Pocket 3 (Supply Chain & Vendor Concentration)

The vendor concentration report shows the top three suppliers accounting for 71% of total COGS. None has been competitively bid since 2022. A 1.5% COGS savings target on the \$90 million COGS base produces \$1.35 million of annual EBITDA. Inventory turns at 4.6 against a peer median of 5.4 imply a 15% inventory reduction opportunity — \$2.9 million of one-time working capital release plus carrying-cost savings of \$175,000 annually. Energy costs flat year over year, no red flag.

Pocket 3 dollar quantification cell: \$1.5 million annual recurring + \$2.9 million one-time WC release. Root-cause hypothesis: stock layers from 2021 supplier disruption never reset; vendors not rebid in three or more years. Candidate workflow: demand forecasting with autonomous reordering, paired with a competitive rebid of the top three suppliers. Implementation timeline: 12 to 16 weeks to Gate 1. Pilot budget: \$150,000 to \$250,000. Payback: 9 to 18 months. Strongest exit-narrative pocket.

Phase 3 — Triage: Pick One

Three pockets sized. Total annual recurring EBITDA opportunity: \$4.8 million. Total one-time working capital release: \$7.0 million. Combined, the diagnostic identifies \$11.8 million of recoverable value against an \$18 million EBITDA base, of which \$4.8 million is annual recurring — within the \$5.3 million canonical figure that anchors this book.

Apply the triage decision rule. The largest dollar opportunity is Pocket 1 at \$2.1 million. The lowest implementation complexity is also Pocket 1 — AP automation is the most mature workflow, the cleanest baseline data, the largest vendor ecosystem. Pocket 1 is the pilot.

Pilot Selection Box, completed:

Selected pilot workflow	Predictive AP automation (50,000 invoices/year)
Pilot owner (named)	[CFO]
Year 1 gross EBITDA target	\$135,000 (pilot scope) / \$1.2M (full-scale projection)

Pilot implementation budget	\$150,000 (one-time) + \$80,000 (annual licence)
Vendor demos scheduled by	Friday of week 2
Pilot kickoff target date	Monday of week 4

That is the worksheet. Twelve cells filled in, plus the pilot selection box. One page. Ninety minutes. The CFO walks the page to the CEO on Tuesday afternoon. The CEO either approves it on the spot or schedules the IC stress test for Thursday. Either way, the next 14 days are described in Chapter 9.

If your worksheet looks materially smaller than this on a \$150 million revenue base — fewer than three pockets, or recurring EBITDA recovery under \$2 million — the most likely explanation is that one of the three input documents was not actually pulled in detail. The second most likely explanation is that you are in the top decile of mid-market operational efficiency, in which case the 90-day refund window is the appropriate response. There is no third explanation.

Chapter 8 — The Three Stress Tests Every IC Will Run

The diagnostic is done. The calculator is filled out. The pilot is named. You take it to the board, or to the IC if you are a PE operating partner, and you present it. You will encounter exactly three objections. They are the same three objections in every IC meeting in every PE firm in 2026, and the responses to them are the same. This chapter is the script. Memorise it, or better, internalise the underlying logic, because the questions will arrive in slightly different language.

Stress Test 1 — "I have seen this before"

The objection in long form: Every consulting firm in the world has been promising EBITDA improvement from technology investments for thirty years. Every firm has its acronym. Every firm has its case studies. The case studies are always the same three companies, presented at every conference, recycled across every deck. We have spent millions on these initiatives in the past and the results have been ambiguous at best. Why is this one different?

The response has three parts.

First, deterministic unit economics. Every previous wave of operational improvement consulting was sold on a vague promise of "efficiency" or "productivity gains." The numbers in the IC memo were qualitative. The post-mortem two years later was inconclusive because the baseline metrics had not been measured rigorously. SEAS pilots are sold on cost per invoice, DSO days, energy cost per square foot — numbers with one decimal place, measured before and after, audit-trail documented. The post-mortem is not inconclusive. Either the cost per invoice dropped from \$4.50 to \$1.80 or it did not.

Second, the implementation cost is one to two orders of magnitude lower than previous waves. A traditional ERP-anchored operational improvement initiative cost \$2 million to \$5 million in software, \$1 million to \$3 million in implementation services, and 18 to 24 months of organisational disruption. SEAS pilots cost \$50,000 to \$150,000 in total commitment. If the pilot fails, you lose the price of an executive's annual bonus. If it succeeds, you scale.

Third, the technology has actually shipped. Previous waves of operational AI required custom development, in-house data science teams, and proprietary integrations. The current generation of agentic AI is shipped product, with documented mid-market deployments, transparent pricing, and standard integrations to NetSuite, SAP, Oracle, and Microsoft Dynamics. The risk profile of a 2026 AP automation pilot is closer to the risk profile of buying Salesforce than the risk profile of building a custom data science capability from scratch.

Stress Test 2 — "Show me the first dollar"

The objection in long form: Pilot economics look fine on a slide. When does this actually generate cash?

Working capital release from AR DSO improvement begins in Week 6 of the AP/AR pilot. A 10-day DSO improvement on \$150 million in revenue is a \$4.1 million one-time working capital release. The release lands in the bank account before the pilot has even reached Gate 1. It is the fastest cash impact in the SEAS framework.

Late-payment penalty elimination begins in Week 4. Direct cost-per-transaction savings begin in Week 8, at Gate 1. By Week 24, at Gate 2, the cumulative pilot savings should be in the \$200,000 to \$500,000 range. There is no capex spike. Agent deployment is operating expense, expensed against the period. The board sees the cost on the P&L in the same period it sees the offsetting savings.

Stress Test 3 — "What if the diagnostic finds nothing"

The objection in long form: We are reasonably well-run. What if we run FLOAT and it finds nothing?

The response rests on the published benchmark record rather than on private claims. Against the APQC, Bain, FTI Consulting, and KPMG benchmark gaps that FLOAT tests for, a mid-market firm that has not run a structured margin-recovery programme will surface at least three quantifiable EBITDA opportunities in the overwhelming majority of cases. The exceptions are companies that are already top-decile AI-native, or companies whose data infrastructure is so degraded that the diagnostic cannot be run on the trailing twelve-month P&L without pre-work. Both outcomes are valuable information. The reason is structural. Mid-market companies that have not specifically addressed each of the three pockets — through a structured programme with named owners, dollar targets, and quarterly tracking — have the leakage. The leakage is not a function of bad management. It is a function of the absence of a specific intervention. Nobody on the executive team is full-time on margin recovery, and so the margin leaks at the rate of 3 to 5% of revenue annually, indefinitely, until somebody runs the diagnostic and names it.

FLOAT has three acceptable outcomes. First, it finds three or more EBITDA pockets totalling \$3M to \$8M, in which case you have a pilot to run and SEAS is the next step. Second, it finds fewer than three pockets and confirms that your company sits in the top decile of mid-market operational efficiency — also valuable information, and the 90-day refund applies. Third, the diagnostic does not run cleanly because the data is not available or the executive team has not engaged with the worksheets — in which case the SEAS Playbook page 16 explains the prerequisite failure modes, and the 90-day refund still applies. There is no scenario in which 90 minutes of structured diagnosis was wasted, and there is no scenario in which a buyer is locked in if the diagnostic does not deliver.

Chapter 8.5 — Common Diagnostic Mistakes and How to Avoid Them

FLOAT is a 90-minute diagnostic with twelve cells on a single page. The methodology is simple. The execution is where most first-time runs fail. This chapter catalogues the seven most common mistakes encountered in diagnostic runs, and the fix for each. Read this before you run your first FLOAT. Read it again after you finish your first FLOAT and before you run your second.

Mistake 1 — Filling cells with adjectives

The most common failure mode. The CFO writes "significant SG&A overhead" or "material pricing exposure" or "meaningful supplier concentration" into the dollar quantification cells. The worksheet is now useless. Adjectives do not get acted on. Adjectives do not survive an IC stress test. Adjectives do not produce pilot budgets.

The fix is mechanical. If a cell does not contain a dollar number with at most one decimal place, the cell is not filled. If the data is not available to produce a dollar number, that itself is the finding — write "data gap: [specific gap]" in the cell, treat closing the gap as a precondition, and move on. The worksheet must produce numbers.

Mistake 2 — Conflating annual EBITDA with one-time working capital release

Pocket 1's DSO improvement and Pocket 3's inventory reduction both produce one-time working capital releases that show on the cash-flow statement, not the P&L. Adding these to annual recurring EBITDA on the worksheet is the fastest way to lose credibility with a finance audience. The board CFO will spot it immediately. The IC will spot it within thirty seconds.

The fix is the dollar quantification cell format used throughout this book: \$X annual recurring + \$Y one-time WC release. Two numbers. Separated. Both important, neither interchangeable. The IC memo presents both, with the recurring number leading because it is the figure that compounds into multiple expansion at exit.

Mistake 3 — Picking three pilots instead of one

The CFO completes the diagnostic, finds three pockets, and recommends pursuing all three in parallel. This is the fastest way to join the 65 to 80% of AI deployments that fail. Three concurrent pilots dilute executive attention, fragment the implementation team, and produce no validated capability across any single workflow.

The fix is the discipline of starting with one pilot, validating it through Gate 1 at week 8 against the ten gate criteria, and only then scaling to the second pocket. The other two pockets go on the Phase 2 backlog with named owners and target start dates, but they do not start until the first pilot is at production accuracy. The discipline feels conservative. It is not. It is the discipline that distinguishes the 20 to 35% of AI deployments that succeed from the 65 to 80% that do not.

Mistake 4 — Running the diagnostic with the wrong people in the room

FLOAT is a CFO-led diagnostic, not a CIO-led diagnostic and not a strategy-team-led diagnostic. The reason is structural. The CFO owns the P&L. The CFO is the only executive whose performance is measured directly in EBITDA. If the CFO is not in the room running the diagnostic, the dollar quantification will be optimistic and the root-cause hypotheses will be soft.

The fix is to put the CFO at the table for the entire 90 minutes, with the controller alongside for the SG&A pocket and the procurement lead alongside for Pocket 3. The CIO, the CRO, and the chief strategy officer can join the read-out at minute 95. They cannot run the diagnostic. They are not measured on the metric the diagnostic produces.

Mistake 5 — Letting the diagnostic drift past 90 minutes

Ninety minutes is not a target. It is the constraint that produces the discipline. A diagnostic that runs to 180 minutes is no longer a diagnostic. It is a workshop. Workshops produce documents nobody reads. Diagnostics produce single pages CFOs walk to CEOs.

The fix is a hard 90-minute timer, visible to everyone in the room. Phase 1 ends at minute 30. Phase 2 ends at minute 60. Phase 3 ends at minute 90. If a phase is incomplete at the deadline, write "data gap: [specific gap]" in the unfilled cells and move on. The worksheet is the deliverable. Completeness of the worksheet matters more than thoroughness of any individual pocket.

Mistake 6 — Skipping the root-cause cell

The root-cause hypothesis cell is the cell most often left blank or filled with vague phrasing. The CFO sizes the dollar opportunity, names the workflow, and skips the root cause because the root cause feels less actionable. This is wrong. The root cause is what the IC challenges in stress test 1. Without a defensible root-cause hypothesis, the recommendation is one consultant deck away from being dismissed as another wave of operational improvement theatre.

The fix is to write the root cause in plain language, in one sentence, naming the structural reason the leakage exists in this specific firm. Not "inefficient processes." Instead: "AP team grew linearly with revenue because every additional million in revenue produced additional invoices, and no automation initiative was funded between 2018 and 2024." That sentence survives the IC. The vague version does not.

Mistake 7 — Picking the pocket with the largest dollar number rather than the right pocket

The triage decision rule in this book is not "pick the largest pocket." It is "pick the largest pocket with the lowest implementation complexity." These are not the same thing. Pocket 3 is often the largest dollar opportunity, but its 9 to 18-month payback and its multi-system integration profile make it a poor first pilot for a firm that has not yet validated any AI workflow. Pocket 1, with its 3 to 6-month payback and its mature vendor ecosystem, is almost always the correct first pilot — even when its dollar value is smaller.

The fix is to run the ROI calculator from Chapter 7 against each of the three candidate workflows and pick the workflow with the best risk-adjusted Year 1 economics, not the largest gross opportunity. The exception is the firm preparing for a sale or refinancing in 18 to 24 months, where Pocket 3's exit narrative dominates and the longer payback is acceptable because the data-room story matters more than the Year 1 EBITDA.

Avoid these seven mistakes and the diagnostic produces a defensible single-page output in 90 minutes. Make any one of them and the diagnostic produces a document that ends up filed in a shared drive, marked "Q4 priority," and never opened again.

Chapter 9 — From Diagnostic to Pilot in 14 Days

The diagnostic is the start. The pilot is the result. The fourteen days between them are the most important fourteen days in the entire SEAS implementation, because this is where most companies lose momentum. The diagnostic gets done on Monday. By Friday, the operating committee has moved on to other priorities. The named pilot becomes a line item on a strategic initiatives list. Two months later, the FLOAT worksheet is in a shared drive folder, marked "Q4 priority," and the company has lost 60 days of EBITDA recovery to inertia.

Sixty days of inertia at \$12,500 per day is \$750,000 gone. That is the price of reading this book and not acting on it.

This chapter is the playbook for the fourteen days. Print it. Tape it to the wall. Follow it.

Day 1 — Monday: The Diagnostic

Block 90 minutes on the calendar. Pull the three documents. Sit with the FLOAT worksheet from Chapter 3 and run the three pockets in order. At the end of 90 minutes, you have a single-page output: three pockets sized in dollars, root causes named, candidate workflows identified. In the next 30 minutes, run the ROI calculator on the candidate workflow from each pocket. Compare. Pick the pilot. Write its name on a separate page in capital letters at the top: PILOT WORKFLOW. Put your name underneath: PILOT OWNER. This is your accountability document. You are now the named owner.

Days 2 to 3 — Tuesday and Wednesday: Vendor Outreach

Identify three vendors that can deliver the named pilot workflow. Email each vendor on Tuesday morning. The email is short. Five sentences. Subject line: "Pilot opportunity — [your company] — [workflow] — Q1 demo." Body: We are a \$X million revenue mid-market company in [industry]. We are evaluating an 8-week pilot of [workflow] beginning [date]. We need a 30-minute demo focused on implementation timeline, pilot pricing, accuracy benchmarks at week 8, and reference customers in our industry. Available [day] or [day] this week. CFO will join the demo.

Three vendors. Three demos. Three responses by Wednesday end of day. If a vendor does not respond within 48 hours, replace them with the next vendor on the list. Vendor responsiveness in the sales cycle is the best available leading indicator of vendor responsiveness in the implementation.

Days 4 to 8 — Thursday through next Monday: Vendor Demos

Schedule the three demos for Thursday, Friday, and the following Monday. Same questions for each. Same scoring sheet. The full SEAS Vendor Matrix — a live Excel workbook scoring 10 agentic AI platforms across 6 weighted criteria — gives you the complete scoring framework. For this 14-day sprint, the seven non-negotiable questions are:

- Implementation timeline to pilot Gate 1 in weeks. Anything over 12 weeks is a red flag for an 8-week pilot.
- Pilot pricing — implementation cost and licensing. Anything over \$200K total for the pilot phase is a red flag.
- Documented accuracy at week 8 for similar deployments. Anything below 90% is a red flag.
- Three reference customers in your industry willing to take a 15-minute reference call. Inability to provide references is the most serious red flag of all.
- Integration to your existing ERP. Pre-built connector or custom development.
- Kill-switch protocol. SEAS standard is 5 minutes. Anything over 30 minutes is a red flag.
- SLA terms with refund triggers. The full SEAS package includes the vendor SLA addendum template with tiered refund triggers (100%/50%/25%) tied to accuracy thresholds.

Days 9 to 14 — Reference Calls, Decision, and Board Approval

Days 9 to 11: Take the three reference calls. Each is 15 minutes. Three questions: How long did your pilot take from kickoff to Gate 1? What was the actual accuracy at Gate 1 versus what the vendor promised? If you were doing it again, would you choose the same vendor?

Days 12 to 14: Draft the board memo. Five sections. Diagnostic results. Selected pilot. Risk mitigation. Decision required. Timeline. Walk the memo to the CEO. Fifteen minutes. The CEO either approves it on the spot or asks for a board vote. Either way, the pilot has a start date by Friday afternoon, and the start date is no more than two weeks out.

That is the playbook. Diagnostic on Day 1. Pilot kickoff scheduled by Day 14, with start date no more than two weeks out. No exceptions, no excuses. The discipline of the fourteen-day window is what separates the companies that capture the EBITDA from the companies that file the FLOAT worksheet in a folder and forget about it. Every

company in the second group is burning \$12,500 a day. Every company in the first group is compounding it into equity value.

From Worksheet to Board Memo: A Template

The FLOAT worksheet is the diagnostic output. The board memo is what gets the pilot funded. The two are different documents. The worksheet is for the CFO who ran the diagnostic. The board memo is for the audience that did not run it and needs to approve the budget for the pilot.

Translating the worksheet into a board memo is the single most under-discussed step in the FLOAT-to-pilot path. Most CFOs assume the translation is straightforward — copy the numbers from the worksheet into a slide deck, add a recommendation slide, send it. This produces decks that get pushed to next quarter for follow-up questions. The decks that get approved on the same Tuesday they are presented follow a tighter five-section structure that this chapter walks through.

Section 1 — The single-sentence ask

The first sentence of the memo is the ask. Not the context. Not the diagnostic. The ask. "This memo requests board approval for a \$150,000 pilot of predictive AP automation, targeting \$1.2 million of annualised EBITDA at full-scale deployment, with a Gate 1 decision point at week 8."

Every word of that sentence is doing work. The dollar amount of the pilot. The dollar amount of the prize. The decision point. The board has, at most, ninety seconds of attention before turning to the next memo. If the ask is not in the first sentence, the ask does not get heard.

Section 2 — Diagnostic results

Three sentences. The total recoverable EBITDA identified by the FLOAT diagnostic. The three pockets named, sized in dollars. The pocket selected as the candidate pilot, with the reason. "FLOAT identified \$4.8 million of annual recoverable EBITDA across three pockets — SG&A efficiency (\$2.1M), pricing drift (\$1.2M), and supply chain (\$1.5M). The pilot recommendation is SG&A efficiency, on the basis of largest dollar opportunity at lowest implementation complexity."

No tables yet. No charts yet. Three sentences. The board now knows the size of the prize and the rationale for the pilot choice.

Section 3 — Pilot economics

This is where the ROI calculator from Chapter 7 lands, condensed to a six-line summary table:

Year 1 implementation	\$150,000 one-time
------------------------------	--------------------

cost	
Year 1 licensing cost	\$80,000 annual
Year 1 gross EBITDA at pilot scope	\$135,000
Year 1 net EBITDA at pilot scope	(\$95,000) — investment year
Full-scale annualised EBITDA (post-pilot)	\$1.2 million
Payback period (full-scale)	4 months

The Year 1 net negative is unfamiliar territory for some board audiences. Address it directly in one sentence: "Year 1 net is negative by design — the pilot scope is intentionally below full enterprise volume to validate the technology before scaling. Year 2 economics, with no further implementation cost, turn solidly positive at the pilot scope alone."

Section 4 — Risk mitigation

The board will ask about risk. Pre-empt the question. Three sentences. The Gate 1 criteria. The kill-switch protocol. The vendor SLA refund triggers. "Pilot is gated at week 8 against ten pre-defined criteria including 90% accuracy, 25% cycle-time reduction, and 95% uptime. Kill-switch protocol returns the workflow to manual processing within five minutes if accuracy drops below threshold. Vendor SLA includes tiered refund triggers (100%/50%/25% of annual licensing cost) if accuracy thresholds are missed at week 8 or 16."

The board now knows the downside is bounded. The risk-mitigation section is the section that converts a sceptical board to a supportive board. Without it, the memo reads as a budget request. With it, the memo reads as a structured experiment with engineered downside protection.

Section 5 — Decision required and timeline

The closing section repeats the ask, names the next steps, and states the timeline. "Decision requested at this meeting: approval of the \$150,000 pilot budget. Vendor demos scheduled for week 1 of next month. Vendor selection by week 2. Pilot kickoff by week 4. Gate 1 review eight weeks after kickoff, within the same quarter."

The board approves the memo or it does not. There is no in-between. The five-section structure forces the board to confront a binary decision rather than to defer to follow-up questions. Most boards, presented with a defensible \$150K experiment with \$1.2M of upside and an 8-week validation gate, approve. The boards that do not approve raise specific questions that this five-section structure makes easy to answer.

What the memo is not

The board memo is not a strategy document. It is not an industry benchmark report. It is not a recap of the AI literature. It is not an explanation of how agentic AI works. The board memo is a one-page recommendation that funds a specific pilot. Anything that does not contribute to the funding decision is removed. The CFO who tries to educate the board on AI fundamentals in the same memo as the pilot ask gets neither outcome.

The full SEAS package includes the board deck template — a 10-slide presentation that maps to this five-section memo structure, plus a 4-slide fund-level deck for PE operating partners reporting up to LPs. Both are pre-formatted and structured around the same five-section logic described in this chapter. The templates are in the SEAS Templates Library.

The CFO who walks this five-section memo to the CEO on the Tuesday after the diagnostic schedules a board vote within two weeks. The CFO who tries to write the comprehensive AI-strategy document instead schedules another diagnostic conversation in Q3.

Chapter 10 — What Comes After FLOAT: The SEAS

This is the chapter where I tell you what I am selling, and why you might want to buy it. I am going to do it directly, because the alternative is the kind of indirect product positioning that wastes your time and insults your intelligence.

The reason the diagnostic and the implementation system are separate documents is that they answer separate questions. FLOAT answers "is there a number, and if so how big, and where." SEAS answers "given the number, what specifically do I do for the next eighteen months to capture it without joining the 65–80% of AI projects that fail." These are not the same problem. Identifying \$5.3 million of leakage is a financial-forensics exercise that takes 90 minutes. Capturing \$4–\$6 million of that leakage cumulatively over 18 months — and producing the audit trail that justifies the multiple expansion at exit — is a programme-management exercise that takes a 309-page playbook, 35 agent specifications, 9 financial models, 20 legal templates, and a CFO willing to sponsor the work for 18 months. Phase 5 of the programme, in months 19 to 24, layers bonus systems on top to lift the cumulative captured EBITDA to \$6 million to \$10 million. None of that is in this book, because none of it is needed until after the diagnostic confirms the number is there.

FLOAT is the diagnostic. SEAS — the Strategic EBITDA Acceleration System — is the implementation system that takes you from the named pilot through the 18-month deployment, the five gates, the governance framework, the legal templates, the financial models, and the exit documentation.

The two are not interchangeable. FLOAT in 90 minutes tells you whether you have the leakage and where it is. SEAS deploys the leakage-capture agents over an 18-month programme (Phases 0–4) that delivers \$4M–\$6M of cumulative captured EBITDA and a \$4M–\$5.5M annual run-rate by month 18. Phase 5 (months 19–24+) layers bonus systems on top — sales, customer service, compliance, ecosystem effects — to lift the cumulative to \$6M–\$10M and the annual run-rate to \$6M–\$10M by month 24+. The audit trail this produces is what justifies the multiple expansion at exit.

What is in the SEAS package

SEAS is delivered as a single package for \$2,997. The package contains five components, each of which would cost more than the entire package price if purchased individually as a consulting deliverable from a major firm.

Component	What It Is	Standalone Cost Equivalent
The SEAS Playbook (Vol. 1)	309-page implementation playbook with 5-phase deployment, financial models,	\$200K–\$500K (McKinsey)

	NIST AI RMF governance framework, and the full 18-month execution plan	diagnostic)
The RAPID SEAS Playbook	26+ page daily execution manual for operating partners and CFOs in the field. Self-contained 12-week toolkit, FLOAT diagnostic, ROI calculator, vendor matrix, board deck template	\$75K-\$200K (Big 4 readiness assessment)
Strategic Agents' Library (Vol. 2)	146+ pages of agent specifications: 9 cross-functional financial engineering agents plus 26 industry-specific agents across 13 verticals, with data inputs, agent logic, output dashboards, vendor options, and governance controls	\$150K-\$400K (custom agent design)
Agent Prompts' Library (Vol. 3)	82+ pages of production-ready prompts for the strategic agents, with the two-pass audit methodology and validation protocol. Copy, paste into Claude or GPT-4 with your CSV data, get directional analysis in 15 minutes.	Included only with SEAS
Templates' Library (Vol. 4)	196+ pages of legal and operational templates: AI Acceptable Use Policy, Vendor SLA Addendum with refund triggers, NHI Access Management Policy, CEO Alignment Charter, IC reporting template, board deck template, Vendor RFP template	\$10K-\$50K (legal counsel drafting)
Project Manager's Runbook	Printable, day-by-day execution reference that carries the operating partner or CFO through the 18-month deployment, the five gates, and the weekly reporting cadence.	Included only with SEAS
Comparative Analysis	Side-by-side benchmark of the SEAS approach against McKinsey, BCG, and Deloitte engagement models on cost, speed, and transferability.	Included only with SEAS
9 Excel Workbooks	1,711 live formulas across Data Fitness Scorecard, PE EBITDA Calculator, MOIC Calculator, ROI Calculator, Vendor Matrix, Implementation Budget, Gantt Template, and more	\$150K-\$400K (bespoke financial modelling)
Board & Fund-Level Decks + Software	10-slide Board Deck, 4-slide Fund-Level Deck, and the EBITDA Pathfinder Pro software application for Windows	Included only with SEAS

The total package price is \$2,997. The total equivalent cost of the standalone components, sourced from major consulting firms or specialised legal counsel, is between \$585,000 and \$1,550,000.

Read that sentence one more time. The price differential is not 10x. It is not 50x. It is roughly 200x to 500x.

The price differential is not because SEAS is lower quality. The price differential is because SEAS is a productised consulting deliverable rather than a bespoke consulting engagement, and the productization removes 99.5% of the cost while preserving the substance.

That is the trade-off. You do not get a McKinsey logo on your slide deck. You do not get a partner from a Big 4 firm taking your CFO out to lunch every other Thursday. You do not get the institutional credibility of a \$400,000 invoice that the board respects because the price was so high it must have been valuable. You get the substance — the financial models, the templates, the implementation plans, the legal documents — at the price of a mid-range office laptop.

If you can run the implementation with your own team, you save the consulting fee. If you cannot, you bring in a consulting firm with the SEAS Playbook in hand and the consulting engagement is shorter, cheaper, and more focused, because the consultant is implementing a defined system rather than designing one from scratch.

The 90-day decision-clarity guarantee

This is the part most buyers do not believe until they read it twice.

Your purchase carries a 90-day, 100% refund guarantee — no questions asked, matched to the full 90-day deployment timeline, so the guarantee never expires before the work is done. Run the FLOAT diagnostic on your firm; it takes 90 minutes. If it does not identify at least 3 actionable EBITDA opportunities for your organization — or if you are unsatisfied for any reason at all — request a refund and it will be processed within 48 hours. No interrogation. No friction. No "can we get on a call to understand why."

90-day deployment timeline. 90-day 100% refund guarantee. 90 days of email support. You're covered for the entire run. Every support email gets a reply within two business days.

This guarantee exists for a specific reason: I am confident in the diagnostic. The benchmark gaps it tests for are documented across the published research, and a mid-market firm that has never run a structured margin-recovery programme sits on the wrong side of them almost by definition. I am putting the guarantee in writing because the asymmetry is real — your downside is bounded at \$0 for a full 90 days,

and your upside is bounded at **\$5.3 million** over 18 months. There is no risk-adjusted scenario in which a rational CFO declines this trade.

The only condition attached to the refund: the Executive Honor Code. Upon receiving a refund, the buyer agrees to permanently delete all copies of the SEAS package from their systems. This is standard enterprise software licensing practice, and it is enforced by mutual trust rather than by DRM, because the alternative is the kind of friction that makes buying SEAS feel like dealing with Oracle, which is the last thing I want.

Who SEAS is for, and who it is not for

SEAS is for three buyers, each with a different use case.

PE operating partners running a portfolio of 5 to 15 mid-market companies. The economics here are dominant. A single deployment across one portfolio company captures \$3M to \$5M of recoverable EBITDA on the first attempt and pays back the SEAS investment in 9 hours or less of recovered margin (at \$3M/year that is \$342/hour; at \$5M/year, \$570/hour). The same playbook deploys across the next four portfolio companies at 40 to 60% lower cost per deployment.

Mid-market CFOs at \$50M to \$500M revenue companies running internally. A \$150M revenue company captures \$5.3M of recoverable EBITDA in the first 18 months, pays back the SEAS investment in less than 6 hours of recovered margin (\$5.3M/year is roughly \$605/hour), and embeds the agentic infrastructure that compounds over the next decade.

Mid-market CFOs at companies preparing for a sale or refinancing in the next 18 to 24 months. The audit-ready governance dossier, the documented EBITDA improvement, and the AI-native exit narrative are modeled to justify a 1 to 3 turn multiple expansion at sale. SEAS pays for itself roughly 12,000 to 36,000 times over in this scenario.

SEAS is not for you if:

- Your company is under \$25M in revenue. The agent licensing economics do not work below this scale. Wait until you are bigger.
- Your company is in active turnaround or restructuring. The playbook assumes operational stability and capital availability for a \$50K to \$150K pilot. Stabilise first.
- Your CFO and CEO are unwilling to put their names on the implementation as executive sponsors. Without that sponsorship, no AI deployment succeeds, regardless of the quality of the playbook. Fix the sponsorship problem first.

- You are looking for a framework to read, not a system to execute. SEAS is 800+ pages of execution tooling. If you want a book to read on a plane, buy something else.
- You need a McKinsey logo to defend the decision to your board. SEAS is productised, not institutional. If the logo is the value, hire the firm that comes with the logo.

If none of those five disqualifiers apply to you, the next step is the link at the end of this chapter. If any of them do, close this book and do not buy SEAS. I would rather have you not buy than have you buy and be disappointed, because disappointed buyers become refund requests and warn their peers, and both of those outcomes cost me more than the \$2,997 is worth.

What you should do this week

Three things, in order.

First, run FLOAT. This week. Not next week, not after the budget cycle, not after the next board meeting. Block 90 minutes on Tuesday or Wednesday morning, pull the three documents, and run the diagnostic. The cost is 90 minutes of your time. The output is the worksheet that tells you whether your company has the leakage and where it is.

Second, run the ROI calculator on the highest-value candidate workflow. Another 30 minutes. The output is the year-one and year-two economics of the pilot.

Third, decide whether to buy SEAS. This is the decision that has been waiting for you on the back cover of this book since you started reading. You now have everything you need to make it informed. You know the diagnostic. You know the framework. You know the trade-offs. You know what is in the package and what is not. You know who it is for and who it is not for. You know the refund terms. The decision is yours.

If you take nothing else from this book, take this: the only risk that compounds is inaction. For a \$150M revenue company with 3% addressable EBITDA leakage, every month of delay is roughly \$375,000 in unrealised margin. Six months of waiting is \$2.25M. Twelve months of waiting is \$4.5M. The SEAS investment is \$2,997. The 90-day refund means your downside is zero for the entire deployment window. There is no rational reading of this arithmetic in which waiting is the correct answer.

Your Downside Is \$0 for 90 Days. Your Upside Is \$5.3 Million.

One click. Instant download. 90-day full refund — no questions asked.

▶ GET THE FULL SEAS SYSTEM — \$2,997 ◀

<https://smartagenticsystems.com/>

90-day 100% refund guarantee — no questions asked. If FLOAT doesn't identify at least 3 actionable EBITDA opportunities, that alone is reason enough.

Frequently Asked Questions

The questions below are the ones a CFO or PE operating partner is most likely to ask before running FLOAT or deciding whether to pursue SEAS. The answers are direct and short. Most are also addressed at greater length somewhere else in the book; this FAQ is a quick reference.

On the diagnostic itself

Q: Does FLOAT really take only 90 minutes?

A: Yes, when the three input documents are pulled in advance and the CFO is in the room. If you allow the meeting to drift into a workshop, it will take three hours and produce a worse output. The 90-minute constraint is the discipline that produces the single-page deliverable. Treat it as non-negotiable.

Q: Can a controller or FP&A manager run FLOAT instead of the CFO?

A: They can run the data work that produces the dollar quantifications. They cannot make the triage decision in Phase 3. The triage decision requires CFO authority because it commits the firm to a pilot budget. Run the data prep without the CFO; do not run Phase 3 without the CFO.

Q: What if our company is not in the \$50M to \$500M range?

A: Below \$25 million in revenue, the agent licensing economics do not work. The pilot cost is not proportional to the recoverable EBITDA. Wait until you are larger. Above \$500 million, the diagnostic still applies but the dollar figures scale beyond the canonical \$5.3 million reference, and the deployment programme becomes more complex. The full SEAS package is calibrated for the \$50M to \$500M range; firms above \$500M typically engage a Big 4 implementation partner alongside SEAS, using the playbook to define the engagement scope rather than to execute it directly.

Q: What if the diagnostic finds nothing?

A: Empirically, this is the third-most-common outcome and signals one of three things. Either the firm is in the top decile of mid-market operational efficiency (rare but real, and useful information), or one of the three input documents was not pulled in detail (most common, and remediable), or the data infrastructure is so degraded that the standard reports cannot be produced (a finding in itself, and a precondition to fix before any AI deployment).

On the SEAS relationship

Q: Do I have to buy SEAS to use FLOAT?

A: No. The FLOAT diagnostic in this book is complete and stands on its own. The worksheet you produce in 90 minutes is yours to act on independently. SEAS is the implementation system that runs from a validated pilot through the 18-month deployment — useful if you are deploying yourself, not required if you are engaging a consulting firm or hiring internally.

Q: When does buying SEAS make economic sense?

A: When the diagnostic identifies \$3 million or more of annual recoverable EBITDA, when the CFO and CEO are willing to sponsor an 18-month programme, and when the alternative is engaging a Big 4 firm at \$400,000 to \$1.5 million for an equivalent scope. SEAS at \$2,997 is roughly 0.2% to 0.7% of the consulting alternative for a comparable deliverable.

Q: What is in SEAS that is not in FLOAT?

A: The 309-page SEAS Playbook (the implementation programme), the RAPID SEAS Playbook (the daily execution manual), the Strategic Agents Library (146 pages, 35 agent specifications), the Agent Prompts Library (82 pages, ready-to-use prompts), the Templates Library (196 pages, 20+ legal and ops templates), the comparative analysis against McKinsey, BCG, and Deloitte, the 9 Excel workbooks with 1,711 live formulas, the 9 sample CSV data files, the board and fund-level presentation decks, and the EBITDA Pathfinder Pro Windows software. FLOAT is the diagnostic; SEAS is the implementation system.

On scale and applicability

Q: My company is in [industry]. Does FLOAT still apply?

A: The three pockets — SG&A efficiency, pricing drift, supply chain — show up in proportions that vary by industry but persist across all of them. Industrial and distribution firms have the largest Pocket 3. Software and services firms have the largest Pocket 2. All firms in the SEAS target range have Pocket 1. The full SEAS Strategic Agents Library includes 26 industry-specific agent specifications across 13 verticals, calibrated for these proportional differences. The diagnostic methodology in this book is industry-agnostic; the implementation specifics in SEAS are industry-aware.

Q: Are the canonical numbers (\$150M revenue, \$5.3M leakage) realistic for my firm?

A: The \$150M reference is the median of the SEAS target range, chosen for ease of calculation. The recoverable EBITDA figure scales linearly with revenue: a \$50M firm should expect \$1.5M to \$2M; a \$75M firm \$1.7M to \$2.6M; a \$250M firm \$8M to \$10M; a \$300M firm \$10M to \$11M; a \$500M firm \$15M to \$20M. The proportional methodology — the three pockets, the red-flag thresholds, the 90-minute structure — is identical across the full range. Only the absolute dollar amounts scale.

Q: Will FLOAT work for a non-PE-owned company?

A: Yes. The diagnostic is identical for PE-owned and non-PE-owned firms in the SEAS target range. The differences are in the IC stress tests (Chapter 8) and the exit-narrative framing (most relevant for PE-owned firms heading toward sale). The PE multiple expansion arithmetic in Chapter 1 is most directly applicable to PE operating partners; the EBITDA recovery itself is fully applicable to any mid-market CFO running the firm internally.

On execution and risk

Q: What is the most common reason a FLOAT-driven pilot fails?

A: Scope creep. The pilot succeeds at week 8, the team is excited, leadership wants to capitalise on momentum, three additional workflows are added before the first one stabilises. By month nine, none of the workflows are at production accuracy. The fix is the Gate 1 discipline: no scaling decision until the first pilot has met at least 7 of 10 gate criteria for at least four consecutive weeks. SEAS encodes this as a structural safeguard. Without it, expect failure.

Q: What if my CEO is not engaged on AI?

A: Do not start the pilot. The CEO Alignment Charter in the SEAS Templates Library exists because PE operating partners learned, expensively, that voluntary CEO commitment to an AI deployment evaporates the moment the next quarter's revenue target slips. Compulsory CEO commitment, written into the bonus structure with 20% to 30% of Year 1 bonus tied to AI-driven EBITDA milestones, does not. If the CEO will not sign the charter, the AI deployment will not succeed regardless of the playbook quality. Address the CEO sponsorship question first.

Q: How does the 90-day refund work in practice?

A: Within 90 days of purchase of SEAS — a window matched to the full deployment timeline — request a refund for any reason and it will be processed within 48 hours. No interrogation, no friction. Running the FLOAT diagnostic first is recommended (if it fails to identify at least 3 actionable opportunities, that is reason enough), but the refund is not conditional on it. The only condition is the Executive Honor Code: upon receiving a refund, the buyer permanently deletes all copies of the SEAS package from their systems. This is standard enterprise software licensing practice and is enforced by mutual trust rather than DRM.

On the author and the methodology

Q: Where do the canonical numbers come from?

A: The canonical numbers — \$150M revenue, \$18M EBITDA, \$5.3M leakage, 8x to 10x multiple expansion — are anchored in published research from Bain, McKinsey, BCG, FTI Consulting, KPMG, Accenture, RAND, and Gartner across 2024 to 2026, calibrated against the operating realities of mid-market firms in the SEAS target range. The conservative capture rates used throughout this book sit below the high-performer figures in the published research. The SEAS Playbook contains source verification ratings for each operational lever, allowing buyers to distinguish high-evidence levers from moderate-evidence levers when building IC presentations.

Q: How was this book written?

A: The intellectual property — the SEAS framework, the FLOAT diagnostic methodology, the financial models, the canonical figures, the three-pocket leakage taxonomy, the gate criteria, and the 14-day execution playbook — originated with the author across two years of independent research and synthesis. The book itself was prepared with strictly regulated AI assistance for prose drafting, structural editing, and copy refinement only. Every number, every threshold, every recommendation is the author's own, verified against source research and operational data.

About the Author

Lalit Kumar is the author of the Strategic EBITDA Acceleration System (SEAS), the RAPID SEAS Playbook, and the FLOAT Diagnostic Framework. He has more than two decades of experience at the intersection of technology strategy, financial performance, and operational value creation, with deep expertise in agentic AI deployment, mid-market operations, and PE-driven multiple expansion.

He was recognised as an Iconic Indian Leader 2025 by the Times of India. His academic credentials include a Bachelor of Science in Mathematics, Physics, and Computer Science; a Master of Computer Applications; and a Bachelor of Laws — a combination that informs his emphasis on diligence-grade documentation, regulatory alignment, and technically grounded implementation.

The SEAS framework was developed over extended research and writing period across 2024 and 2025, drawing on published research from McKinsey, Bain, BCG, Deloitte, Accenture, FTI Consulting, KPMG, RAND, and Gartner; on the NIST AI Risk Management Framework; and on the operational realities of mid-market companies in the \$50M to \$500M revenue range. Every financial model, every legal template, and every implementation checklist in the SEAS package has been independently verified for internal consistency, factual accuracy, and source attribution. The SEAS Playbook contains source verification ratings for each operational lever, so that buyers can distinguish high-evidence levers from moderate-evidence levers when building IC presentations.

Your Next Step

You have read the book.

You have the diagnostic.

You know the numbers.

There are only two paths forward from this page.

Path 1 — Close this book.

File it on your Kindle. Promise yourself you will get to it next quarter. Continue the weekly rhythm of operating committee meetings in which someone asks about AI and someone else gestures vaguely at a roadmap. Six months from now, you will have burned \$2.25 million of unrecovered margin. Twelve months from now, \$4.5 million. By the time you get around to running FLOAT, the arbitrage window will have narrowed by a third and the premium exit multiple will be harder to defend.

This is the default path. It is the path 80% of readers will take. It is the path that produces the 65–80% failure statistic. It is also the path that requires zero effort and therefore feels safest in the moment.

It is not safest. It is **the most** expensive thing you can do this year.

Path 2 — Click the link: <https://smartagenticssystem.com/>

Download the full SEAS System. Run FLOAT against your trailing twelve-month P&L before Friday. Identify three EBITDA pockets. Pick one. Schedule the vendor demos the following Tuesday. Walk a board memo to the CEO two weeks after that. Kick off the pilot by the end of the month.

By Week 8, you have a validated pilot. By Week 24, you have \$200,000 to \$500,000 of recovered margin on the P&L. By Week 52, you have \$2.5M to \$4M. By Week 78, you have \$4M to \$6M. By exit, you have the audit-ready governance dossier that is modeled to justify 2 to 3 additional turns of EBITDA multiple.

If any step along the way fails, the 90-day refund window bounds your SEAS downside at zero for the entire deployment period, and the gated pilot structure bounds your vendor-side downside at the \$50,000 to \$150,000 pilot budget. You cannot lose more than the price of one executive bonus. You can gain more than the value of five.

APPENDIX A

The FLOAT Diagnostic Worksheet

90-Minute EBITDA Leakage Diagnostic — Single-Page Output

Context

This worksheet is the deliverable of the FLOAT diagnostic. It is one page. It is filled in 90 minutes. It produces three EBITDA pockets sized in dollars, ranked by speed-to-cash, with one named pilot workflow and a named owner. That single page is what you carry to your CFO, your CEO, or your investment committee.

FLOAT is the diagnosis, not the treatment. The worksheet does not prescribe an 18-month transformation programme. It tells you (a) whether your firm has recoverable EBITDA leakage, (b) where the leakage sits, and (c) which pocket to attack first. The implementation system that follows the diagnostic — vendor selection, pilot deployment, gate criteria, scaled rollout — is the SEAS programme described in Chapters 9 and 10.

Use this worksheet when: (i) you are evaluating an AI-related EBITDA opportunity for a \$50M–\$500M revenue mid-market firm, (ii) you have not run a structured leakage diagnostic in the last 12 months, (iii) you can sit with the three required documents for 90 minutes uninterrupted.

Pre-Flight: The Three Documents

If you cannot pull these three documents in 15 minutes, stop. Your data infrastructure is the leakage, and that is itself a finding. Document the data gap on the worksheet under Pocket 1 and proceed no further until the gap is closed.

#	Document	What it must contain
1	Trailing 12-month P&L	Revenue, COGS, gross profit, SG&A by line item, operating income, EBITDA
2	SG&A breakdown by department	Finance & Accounting, Customer Ops, IT/Helpdesk, Supply Chain Ops, HR, Legal, Sales/Marketing G&A —

#	Document	What it must contain
		each as \$ and % of revenue
3	Vendor concentration report	Top 20 vendors by spend; total annual spend per vendor; date of last competitive bid; auto-renewal status

The 90-Minute Structure

Three phases. 30 minutes each. Two short coffee breaks between phases is reasonable. Less than two is heroic. More than two and you are not running FLOAT — you are running a workshop.

Phase 1 — Find (30 min). Pocket 1, the SG&A Efficiency Gap. Run the five red-flag checks against the SG&A breakdown. Sum the recoverable dollars. Flag working-capital release separately.

Phase 2 — Opportunity (30 min). Pockets 2 and 3 in parallel. Pocket 2 is Pricing and Contract Drift, run against the customer master. Pocket 3 is Supply Chain and Vendor Concentration, run against the vendor concentration report and inventory turns metric.

Phase 3 — Triage (30 min). Rank the three pockets by speed-to-cash and implementation complexity. Pick one as the candidate pilot. Name the workflow, the owner, the Year 1 EBITDA target, and the kickoff date.

Pocket 1 — SG&A Efficiency Gap (Reference Tests)

Pull the SG&A breakdown. Run these five checks in order. Each takes about five minutes. Sum the recurring savings; the WC release is a separate one-time number on the cash-flow statement, not annual EBITDA.

Check	Red flag	Dollar logic
SG&A as % of revenue vs peer median	Exceeds median by ≥ 2 percentage points	2pp gap $\times$ revenue (e.g., 2pp $\times$ \$150M = \$3M)
Finance & Accounting as % of revenue	$> 3\%$ (best-in-class $\approx 1.5\%$)	Gap $\times$ revenue $\times$ $\sim 25\text{--}30\%$ recoverable in 12 months
Cost per invoice processed	$> \$3\text{--}4$ per invoice	Volume $\times$ (\$current – \$1.80 target)
AR DSO vs industry median	> 45 days, no automation	DSO-day cut $\times$ (revenue $\div$ 365) = one-time WC release (not annual EBITDA)
AP/AR FTEs per \$100M revenue	> 8 (benchmark 4–6)	Excess FTEs $\times$ $\sim \$80\text{K}$ loaded cost each

Pocket 2 — Pricing & Contract Drift (Reference Tests)

Pull the customer master. Filter for contracts with last-modified date older than 18 months. Sum the annual contract value (ACV) of that filtered list. Multiply by an assumed 3% recoverable price increase — that is your Pocket 2 floor. The ceiling, in companies with severe pricing neglect, is closer to 10% on the legacy-priced subset.

Check	Red flag	Dollar logic
% of contracts not repriced in 18–24 months	$\geq 10\%$ of revenue on legacy pricing	0.5–3% margin recovery $\times$ revenue base (e.g., 1% $\times$ \$150M = \$1.5M)
SKU margin variance across product lines	$\geq 5\%$ spread, no active management	Bottom-quintile SKUs are repricing targets
Average discount depth vs list price	$> 15\%$ average discount	Discount leakage without enforcement

Check	Red flag	Dollar logic
Auto-renewal contracts without escalation	Renewals lacking CPI / cost-pass-through	Repricing trigger identified at next renewal

Pocket 3 — Supply Chain & Vendor Concentration (Reference Tests)

Pull the vendor concentration report and the inventory turns metric. Pocket 3 is the largest absolute-dollar pocket in industrial, manufacturing, and distribution firms. The payback is longer (9-18 months) but the exit-diligence narrative is the strongest of the three.

Check	Red flag	Dollar logic
Top 3 suppliers as % of total COGS	>60% supplier concentration	1-3% margin opportunity (\$1.5M-\$4.5M on \$150M base)
Last competitive bid for top vendors	No competitive bid in 2+ years	0.5-1% COGS savings recoverable
Inventory turns vs peer median	>15% below peer median	15% inventory cut = one-time WC release $\approx$ 15% of inventory value
Energy cost variance YoY	>10% YoY without volume increase	8-15% energy cost reduction achievable

The Worksheet

Twelve cells. Filled in with specific numbers, not adjectives. At the end of 90 minutes, the grid below is the entire deliverable.

	Pocket 1 — SG&A Efficiency Gap	Pocket 2 — Pricing & Contract Drift	Pocket 3 — Supply Chain & Vendor Concentration
Red-flag check (which thresholds tripped?)	e.g., F&A is 6.8% of revenue; cost per invoice \$4.20; DSO 52 days	e.g., 18% of revenue on contracts >24 months old	e.g., top-3 suppliers = 71% of COGS; no bid since 2022
Dollar quantification (annual recoverable / one-time WC)	\$ ___ M annual recoverable + \$ ___ M one-time WC release	\$ ___ M annual recoverable (100% margin flow-through)	\$ ___ M annual recoverable + \$ ___ M one-time WC release
Root-cause hypothesis (why this leakage exists)	Headcount grew linearly with revenue; no automation of AP/AR	Sales discounts baked into auto-renewing contracts; CRO measures revenue retention, not margin	Stock layers from 2021 disruption never reset; vendors not rebid
Next-action workflow (candidate pilot)	AP automation (8-10 wk deploy, 3-6 mo payback)	Contract Value Capture agent (6-12 mo payback)	Demand forecasting + autonomous reordering (9-18 mo payback)

Triage Decision Rule

Triage decision rule

Pick the pocket with the largest dollar opportunity that has the lowest implementation complexity. For most \$50M-\$500M companies, that is Pocket 1 (AP automation). If the firm is heading toward a sale or refinancing in 18-24 months, choose Pocket 3 (demand forecasting) — the exit-diligence narrative is stronger. If pricing has been neglected for three or more years and a CPQ tool is already in place, choose Pocket 2 — fastest payback, 100% margin flow-through.

Pilot Selection Box

After ranking the three pockets, complete the box below. This is the artifact you walk to the CEO or carry to the IC meeting.

Selected pilot workflow	
Pilot owner (named)	
Year 1 gross EBITDA target	
Pilot implementation budget	
Vendor demos scheduled by	
Pilot kickoff target date	

Print this worksheet. Block 90 minutes on a Tuesday or Wednesday morning. Pull the three documents. Sit with the page. Fill the twelve cells. The result is yours, regardless of what you do next.

Guidelines for Filling the Worksheet

1. Use specific numbers, not adjectives.

"Significant" is not a number. "Material" is not a number. "\$2.1M annual recoverable plus \$4.1M one-time WC release" is a number. If a cell cannot be filled with a number, either the data is not available (a finding in itself) or the check does not apply to your firm (mark it N/A).

2. Working capital release is not annual EBITDA.

Pocket 1's DSO improvement and Pocket 3's inventory reduction produce one-time working capital releases that show on the cash-flow statement, not the P&L. Keep them on a separate line in the dollar quantification cell so the IC sees both numbers without confusion.

3. If you cannot fill a cell in 5 minutes, the data is the finding.

A firm that cannot pull SG&A by department in 5 minutes has a data infrastructure problem. A firm that cannot identify its top-3 suppliers by spend has a procurement governance problem. Document the gap in the cell, name it as the root cause, and treat closing the gap as a precondition before any pilot.

4. The next-action workflow must be a named, vendor-addressable workflow.

"Improve AP" is not a workflow. "Predictive AP automation processing 80% of invoices end-to-end without human touch" is a workflow. The test: can you write a five-sentence vendor outreach email asking for a demo of this workflow? If yes, it is a workflow. If no, refine it.

5. Pick one pilot. Not three.

The discipline that separates the 20-35% of AI deployments that succeed from the 65-80% that fail is the discipline of starting with one pilot, validating it through the Gate-1 criteria, and only then scaling to the second. The worksheet identifies three pockets; the IC memo names one. The other two go on the Phase-2 backlog.

6. The output is one page, not ten.

If the worksheet runs to a second page, the diagnostic has drifted into a workshop. Compress. Use bullet fragments, not sentences. The CFO reading this worksheet has 90 seconds before the next meeting. Respect the constraint.

7. Date the page.

The worksheet is a snapshot of the trailing-twelve-month P&L on a specific day. Write the diagnostic date in the header. Re-run the diagnostic annually.

After the Worksheet

The worksheet is the start, not the end. The 14-day path from a completed worksheet to a kicked-off pilot is described in Chapter 9 of the FLOAT book: vendor outreach Days 2-3, vendor demos Days 4-8, reference calls Days 9-11,

board memo and CEO sign-off Days 12-14. The implementation system that runs from pilot kickoff through the 18-month deployment, the five go/no-go gates, the governance framework, and the exit documentation is the SEAS programme.

— Lalit Kumar

The Strategic EBITDA Acceleration System

\$2,997 — Instant Download — 32 Files — 800+ Pages

<https://smartagenticsystems.com/>

- ✓ **The SEAS Playbook (Vol. 1) — 309-page implementation playbook**
 - ✓ **The RAPID SEAS Playbook — 26-page daily execution manual**
 - ✓ **Project Manager's Runbook — printable execution reference**
 - ✓ **Strategic Agents Library (Vol. 2) — 146 pages, 35 agent specs**
 - ✓ **Agent Prompt Library (Vol. 3) — 82 pages, 35 ready-to-use prompts**
- ✓ **Templates Library (Vol. 4) — 196 pages, 20+ legal and ops templates**
 - ✓ **Comparative Analysis — SEAS vs. McKinsey, BCG, Deloitte**
- ✓ **9 Excel Workbooks — 1,711 live formulas, no broken references or circular logic**
 - ✓ **9 Sample CSV Data Files — populated for immediate testing**
 - ✓ **Board Presentation Deck**
 - ✓ **Fund-Level LP Presentation Deck**
- ✓ **EBITDA Pathfinder Pro — Windows software application(Ready-to-Run)**

► **DOWNLOAD SEAS NOW — \$2,997** ◀

<https://smartagenticsystems.com/>

PE firm portfolio pricing: \$1,997 per license (minimum 3 licenses).

90-day 100% refund guarantee — no questions asked. If FLOAT doesn't identify at least 3 actionable EBITDA opportunities, that alone is reason enough.

In the beginning I asked you before you venture into the SEAS, make sure you can FLOAT.

And you have just learned how.

The only step left is to click the link:

<https://smartagenticsystems.com/>

— Lalit Kumar

© 2026 Lalit Kumar. All rights reserved.