THE LININGER GROUP

PORTCO AI PLAYBOOK SERIES · NO. 3

Operational Efficiency

How to take cost out of the work itself, without adding AI run costs you'll pay forever.

> For sponsors who need margin they can see in the monthly close, and for the leaders who need a way to get there without a two-year program.

$18M
Cost out in Year 2 of a $1.1B integration

24%
Shorter global lead times

90 days
To a unit-cost result Finance will sign

© 2026 The Lininger Group. May be shared with attribution to Rebecca S. Lininger and LiningerGroup.com.

OperationalAlpha@LiningerGroup.com

EXECUTIVE SUMMARY

Adoption isn't efficiency

If you're a sponsor, you've seen the AI licenses on the expense line. You haven't seen the unit cost move.

If you're running the business, your people are using AI. They write emails faster and summarize meetings. But the cost to process an invoice, turn a quote, or close the month hasn't changed, because the process underneath is the same one you had before AI.

Efficiency comes from redesigning the work: strip the steps that don't carry weight, automate the linear logic once, and use AI only where judgment or messy inputs need it. Then measure the result in unit cost, the number Finance already trusts.

WHAT YOU WILL LEAVE WITH

1. **A way to tell real gains from activity**
   A four-level test that separates experiments and adoption from workflow redesign that shows up in the P&L.

2. **A case-validated method**
   Drawn from a $1.1B integration that cut 50+ competing metrics to two and took $18M of cost out, and from agentic AI that added $2M+ a month in quoted volume in its first three months.

3. **A 90-day framework**
   That redesigns one or two workflows and ends with a before-and-after unit cost Finance signs.

> *AI on top of a broken process gives you a faster broken process, and a token bill that grows every month.*

CONTENTS

01	**Diagnosis**	Four reasons efficiency programs stall, and the four-level test
02	**Case studies**	Dresser integration and Ecolab tender automation
03	**The 90-day playbook**	Three phases, each with an output you can check
04	**The numbers**	What each phase takes and Year 1 scenarios
05	**Week 1 without outside help**	Seven actions that need only internal access
06	**Three AI exercises**	Prompts for workflow, metric, and AI spend scans
07	**Appendix**	Diagnostic scorecard and glossary

SECTION 1 OF 7

Diagnosis: why efficiency programs stall

Most companies already have tools, licenses, and a few pilots. The cost structure still looks the same. There are four usual reasons.

1 **Adoption without redesign**

AI was added to the existing process. People save minutes, but the steps, handoffs, and staffing are unchanged, so nothing shows up in unit cost or the hiring plan.

2 **Too many metrics, no owner**

Every function tracks its own numbers. Nobody can say which two measures run the business, so improvement effort spreads thin and no one owns the result.

3 **Manual handoffs and re-keying**

Work moves between people, spreadsheets, and systems by hand. Every handoff adds delay, errors, and a person whose job is to move data rather than make decisions.

4 **Run costs nobody tracks**

AI usage, compute, and licenses grow every month with no link to a workflow or a benefit. Savings get claimed; costs get ignored.

THE FOUR-LEVEL TEST

Each level is useful. Only the top two change the P&L in a way an investor will pay for.

Level	What it looks like	How it's measured
1 · Experiments	Teams try tools; a few pilots	Number of pilots and anecdotes
2 · Adoption	Licenses rolled out; individuals work faster	Usage and self-reported time saved
3 · Workflow redesign	One process rebuilt: automation does the steps, people handle exceptions	Unit cost, cycle time, and error rate for that workflow
4 · Transformation	Several core workflows rebuilt; roles, spans, and hiring plans change	Margin, growth, or capacity at the company level

The quick test

Can you name one workflow where the steps, the staffing mix, or the cycle time changed, and show a before-and-after number Finance agrees with? If not, you have adoption. That's a fine start, but it isn't efficiency yet.

SECTION 2 OF 7

Case studies from the operating record

Both cases come from Rebecca S. Lininger's corporate leadership roles at GE Oil & Gas and Ecolab.

CASE A · GE OIL & GAS: YEAR 2 OF A $1.1B INTEGRATION

50+ → 2	**24%**	**$18M**	**33 sites**
Competing metrics cut to two	Shorter global lead times	Direct savings	In 16 countries

The context. Year two of integrating a $1.1B Dresser business after GE's $3B acquisition: 33 sites in 16 countries, eroding margins, strained channel partners, and 50+ competing corporate metrics. The two-year pro forma was at risk.

What was done. A Master Black Belt reset run at the sites, not from a program office. 50+ metrics were cut to two delivery KPIs, lead time and first pass yield, and every site cadence, escalation path, and leadership review was rebuilt around them. The full prospect-to-cash systems integration was completed across the acquired portfolio.

The result. Margin erosion reversed. Global lead times fell 24%, the two-year pro forma was beaten, and the integration delivered $18M in direct savings.

- Two load-bearing metrics instead of 50+, so every site worked on the same problem.
- Work redesigned at the sites, where the cost lives.
- Systems integrated once, end to end, instead of patched at each handoff.

CASE B · ECOLAB: AGENTIC AI FOR TENDER RESPONSES

$2M+	**20+**	**4 of 5**
Added quoted volume per month, first three months	AI and IoT products shipped	Global business units

The context. EU tenders required detailed ESG responses, assembled by hand for each bid.

What was done. Built agentic AI, using retrieval from approved company material, that automated the ESG tender responses.

The result. $2M+ a month in added quoted volume in the first three months.

What this means for your business

Both cases started with the work, not the tool: fewer measures, fewer handoffs, and automation placed where volume and repetition justify it. The same sequence works in a $50M business. The tools are cheaper now; the method is the same.

SECTION 3 OF 7

The 90-day operational efficiency playbook

Three phases. Each has an output you can check on its own, so your sponsor sees progress before Day 90.

PHASE 1 · WEEKS 1–4 Map the work and pick one or two workflows

Workflow inventory	List your ten largest workflows by cost, headcount, or cycle time: quote to order, invoice to pay, order to ship, month-end close.
Baseline unit cost	For the top candidates, measure cost per transaction, cycle time, and error rate today. Finance validates every baseline.
Pick one or two	Choose workflows big enough to show in the P&L, with enough volume and repetition to automate.
Set the two metrics	Name the two measures that define success and the kill criteria, before any build starts.

PHASE OUTPUT

A one-page brief per workflow: the baseline, the target unit cost, the two metrics, and the kill criteria.

WHAT IT TAKES

Process owners and finance for a two-day working session, plus time to pull cost and volume data.

PHASE 2 · WEEKS 5–8 Redesign and pilot

Strip the steps	Map the current workflow, then design the ideal path: minimum steps, only the constraints that carry weight. Use it as the yardstick.
Automate the linear logic	Route rules-based steps to straight-line automation you pay for once: matching, routing, data entry, standard documents.
AI where judgment is needed	Use AI for messy inputs and exceptions: reading documents, drafting responses, classifying cases. People review what matters.
Pilot on real volume	Run the redesigned workflow on real work and real data for at least two cycles. Track unit cost and exceptions weekly.

PHASE OUTPUT

A redesigned workflow running on real volume, with weekly unit cost, exception rate, and AI run cost per transaction.

WHAT IT TAKES

A process owner, an automation or AI builder, and system access. Effort depends most on your data and integrations, which Phase 1 sizes.

PHASE 3 · WEEKS 9–12 Prove it and scale it

Finance sign-off	Compare unit cost, cycle time, and errors with the baseline. Finance signs the number, net of AI run cost.
Run cost controls	Track AI usage and cost per transaction for each workflow, with a monthly budget and an owner.
Capacity plan	Decide what the freed capacity is for: absorb growth, redeploy people, or change the hiring plan. Decide before you scale.
Next wave	Pick the next workflows from what worked, and write the operating manual for the ones now in production.

PHASE OUTPUT

A signed before-and-after unit cost, a run-cost dashboard, a capacity plan, and the next-wave brief.

WHAT IT TAKES

Finance, the process owner, and HR for the capacity plan, with an owner named for each workflow in production.

SECTION 4 OF 7

The numbers: what to expect and when

These ranges reflect lean, focused implementations. Adjust for your revenue base, cost structure, and data quality.

WHAT EACH PHASE TAKES

Phase	Timing	Primary work
1 · Map	Weeks 1–4	Workflow inventory, baselines, metric and kill criteria
2 · Redesign	Weeks 5–8	Process redesign, automation, AI on exceptions, pilot
3 · Prove	Weeks 9–12	Finance sign-off, run-cost controls, capacity plan

Outside spend depends mostly on your systems and data, so size it in Phase 1. Straight-line automation is paid for once; AI usage is a running cost, so track it from day one.

EFFICIENCY SCENARIOS, YEAR 1

Illustrative, modeled on a $75M revenue business; scale proportionally and replace with your own baselines.

Scenario	What happens	Cost out	Margin
Conservative	One workflow redesigned, partial volume	$0.3–0.6M	+40–80 bps
Base	Two workflows redesigned, full volume	$0.8–1.5M	+100–200 bps
Upside	Three workflows, capacity redeployed to growth	$1.5–3M	+200 bps or more

WHERE THE RETURN COMES FROM

Lever	Typical effect
Fewer steps and handoffs	Shorter cycle times and fewer errors before any technology is added
Straight-line automation	Lower cost per transaction, with no running usage cost
AI on exceptions	People spend their time on judgment calls, not data movement
Run-cost control	AI spend tied to a workflow and a benefit, so net savings hold

The 150 bps threshold

Any efficiency investment that can't defend 150 basis points of margin in Year 1, net of AI run costs, doesn't meet the bar. The base scenario clears it.

SECTION 5 OF 7

Getting started: Week 1 without outside help

Seven actions that need only internal access. Finish them before any vendor conversation.

1 **List your ten largest workflows**
Rank them by cost, headcount, or cycle time. The top three are your candidates.

2 **Time one transaction end to end**
Follow a single invoice, quote, or order through every step. Count the days, the people, and the systems it touches.

3 **Count the handoffs and re-keys**
Every time data moves by hand between people or systems is a candidate for automation.

4 **Inventory your AI tools and licenses**
Include the ones employees bought on their own. Note the monthly cost and who uses each.

5 **Pull three months of AI and software bills**
Tie each line to a workflow. Any cost you can't tie to a workflow is a run cost with no owner.

6 **Cut your metrics to two per workflow**
If a team tracks more than five measures, ask which two would tell you it's working.

7 **Define your one workflow**
Which workflow, redesigned, would move unit cost enough to matter? Write it in one sentence and share it with your sponsor.

COMMON PITFALLS

Pitfall	How to avoid it
Automating the current process	Redesign first. Automating a bad process locks in its cost.
Piloting on demo data	Pilot on real volume and real exceptions, or the result won't hold.
Ignoring run costs	Report savings net of AI usage, compute, and licenses, every month.
Claiming gains without changing plans	If staffing and hiring plans don't change, the savings aren't real yet.

SECTION 6 OF 7

Three AI exercises you can run this week

Each exercise uses tools you can access today, runs on your own material in an afternoon, and maps to one failure mode. They're the manual version of what Phase 2 builds: if any output surprises you, that's your Phase 1 target.

Before you start: never put company data into a free AI tool.

Free versions of AI tools may use your inputs to train their models. A paid Claude subscription starts at about $20 a month; Google Workspace with Gemini starts at about $14 a month. Either gives you business-grade data protection. Confirm with IT or legal before running any exercise below.

1 Workflow redesign scan

Addresses: Adoption without redesign

Find the steps that don't carry weight, and which of the rest need AI versus simple rules.

Set up: Claude (paid) or Gemini (Google Workspace). Paste a written description or SOP of one workflow, step by step.

PASTE THIS PROMPT

I'm sharing the step-by-step process for our [workflow, such as invoice processing]. Each step lists who does it, which system they use, and how long it takes. Tell me: (1) which steps add no value for the customer or the business and could be removed, (2) which remaining steps follow fixed rules and could be automated without AI, (3) which steps need judgment or read messy inputs, where AI could help, and (4) the redesigned workflow with the fewest steps. Put the result in a table I can share with my CFO.

What you'll get	A redesigned workflow and a clear split between rules-based automation and AI, the starting point for Phase 2.
Time required	45–60 minutes, including prep

2 Metric consolidation

Addresses: Too many metrics, no owner

Cut a long KPI list to the two measures that tell you whether a workflow is working.

Set up: Same tool. Paste the metrics your teams report today for one area, with a one-line definition of each.

> **PASTE THIS PROMPT**
>
> Here are the metrics our [operations or finance] teams report each month, with definitions. Tell me: (1) which metrics measure the same thing in different ways, (2) which are activity measures rather than outcome measures, (3) the two metrics that best show whether this area is improving cost, speed, or quality, and (4) how often each of the two should be reviewed and by whom.

What you'll get	A two-metric scorecard and a list of measures you can stop reporting.
Time required	30–45 minutes, including prep

3 AI spend scan

Addresses: Run costs nobody tracks

Tie your AI and software spend to the workflows it's meant to improve.

Set up: Same tool. Export three months of AI and software invoices or a license list, with the team that uses each. Leave out personal names.

> **PASTE THIS PROMPT**
>
> I'm sharing three months of AI and software spend: tool, monthly cost, number of users, and the team using it. Tell me: (1) which tools overlap in what they do, (2) the monthly cost per active user for each, (3) which spend I can't tie to a specific workflow or result, and (4) the three changes that would cut spend the most without hurting work that's in production.

What you'll get	A spend map by workflow, a list of overlaps, and a short list of cuts.
Time required	30–45 minutes, including prep

SECTION 7 OF 7

Appendix: diagnostic scorecard and glossary

Score each statement from 1 (strongly disagree) to 5 (strongly agree). Your highest-scoring cluster is your Phase 1 target.

Score (1–5)	Statement
ADOPTION WITHOUT REDESIGN	
___	We've added AI tools, but our unit costs haven't changed.
___	We can't name a workflow whose steps changed because of AI.
___	Our AI results are measured by usage or surveys.
TOO MANY METRICS	
___	Most teams report more than five metrics.
___	Leaders disagree on which numbers matter most.
___	No one person owns the result for our largest workflows.
HANDOFFS AND RE-KEYING	
___	Data is re-entered by hand between systems.
___	Our largest workflows pass through more than three people or systems.
___	Cycle times vary widely and we can't explain why.
RUN COSTS	
___	We can't say what we spend on AI each month by workflow.
___	Employees use AI tools we didn't approve.
___	Our savings reports don't subtract AI costs.

GLOSSARY

Unit cost

The cost to complete one transaction in a workflow, such as one invoice, quote, or order. The measure Finance can verify.

Straight-line automation

Rules-based automation for steps that always follow the same logic. Paid for once, with no usage cost.

Choreographed Intelligence™

The Lininger Group's Lean process architecture: complex exceptions go to AI, and linear business logic goes to automation you pay for once.

AI$ Gauge™

The Lininger Group's framework for showing AI cost in dollars by person, query, and workflow, so run costs stay visible.

150 bps

150 basis points, or 1.5 percentage points of margin, net of AI run costs. The bar for a defensible efficiency investment.

Want help putting this to work?

The Lininger Group leads the workflow and unit-cost work alongside your team on a monthly retainer. The first 30 days set the plan and the baselines.

OperationalAlpha@LiningerGroup.com · LiningerGroup.com