

The digital supply chain

Why the cost of information is becoming the single most important line to optimize in a 50-to-200-employee business — and how to measure it.

For owners and executives of companies with 50 to 200 employees.
Manufacturing, distribution, professional services, construction,
transportation, food processing.

THE THESIS IN ONE SENTENCE

Artificial intelligence is not a new expense line added to your technology budget. It is a rework of the cost of every one of your departments. The executive who treats it as a subscription keeps their cost structure. The one who treats it as a supply chain recovers points of margin that no subscription will let competitors catch up on.

CONTENTS

1. Your company has two supply chains
2. What you already pay for, and don't count
3. The framework: four steps, one chain
4. The invisible step: the cost of a decision
5. Mapping your chain like an org chart
6. Six sectors: today and tomorrow
7. The only two equations that matter
8. Why the gap doesn't close
9. The most common blockers
10. Where to start, and in what order
11. Where Agent Libre fits

A NOTE ON THE NUMBERS

The amounts in this document illustrate a method; they are not a market benchmark. They are built on a reference company described in chapter 2 and on fully loaded hourly rates you can replace with your own. External sources are listed in the appendix; those that come from vendors or extrapolations are flagged as such. A number you have not recomputed with your own data should not be the basis of a business decision.

Your company has two supply chains

The first one you know. Materials come in; they are transported, stored, transformed, delivered. Every step has a supplier, a lead time, a unit cost. When margins tighten, you walk down that chain step by step, and you generally know where to look.

The second one moves information instead of materials. A request for pricing comes in and becomes a quote. An order comes in and becomes a work order, an invoice, a confirmed delivery. A site report becomes a change order — or doesn't. This chain runs through every department and employs a significant share of your people. It does have a few known suppliers — the ERP or CRM vendor, the managed IT provider, the web agency — but each covers only a segment, and nobody is accountable for the whole. It has no unit cost, no owner, and it appears nowhere in your financial statements under that name.

This document is about that second chain. It argues that its cost is now the single most important line a 50-to-200-employee company has to optimize, that it can be measured with the means of a company that size, and that the company that measures it first gains a cost-structure advantage its competitors will not easily close.

27 FTEs

EMPLOYEES WHOSE JOB IS MOVING
INFORMATION, OUT OF 120

\$2.75M

ANNUAL COST OF THIS CHAIN IN THE
REFERENCE COMPANY, OR 12.5% OF
REVENUE

< 1 / 5

THE SHARE OF THAT COST VISIBLE IN
THE TECHNOLOGY BUDGET; THE REST
IS SCATTERED ACROSS PAYROLL

Why now

This second chain has always existed. What is changing is that the mechanical work present at every one of its steps — entering an order, sorting emails, filtering requests, finding a file, reading a drawing, reviewing a first draft — has just moved from the cost of an employee-hour to the cost of compute, billed by usage. AI is not a step in the chain: it is an expense that shows up at every step, wherever a human was doing repetitive work.

In conversations with executives, this shift is generally well understood, and often reduced to a number: a few hundred dollars a month in subscriptions, today and for the decade ahead. That number correctly describes current usage: a few individual licences, used at the edges of the work. It does not describe what happens when AI actually takes over part of the mechanical work of steps 1, 2 and 4. A company that hands data entry, information retrieval and decision preparation to automated systems consumes processing continuously, and the bill follows: at current public rates from model providers, sustained usage runs between \$5,000 and \$10,000 per month per equipped employee — the order of magnitude of a salary, not a subscription.

That is not bad news: this cost does not add, it replaces. It absorbs hours that were paid as salaries in steps 1, 2 and 4, at a lower unit price. But it makes the AI layer the line whose unit cost matters most — the one to negotiate, measure and optimize the way you do raw materials.

In other words, this shift does not create a new expense line. It moves a cost that already existed, scattered across payroll and managed by no one because it was neither visible nor reducible, into a line that can be measured and negotiated.

What this document builds

Chapter 2 establishes what the chain costs today in a reference company, line by line. Chapter 3 gives its structure in four steps. Chapter 4 addresses the most expensive and least measured step, coordination, and proposes a method to price a decision. Chapter 5 explains how to map the chain the way you draw an org chart. Chapter 6 applies the framework to six sectors. Chapters 7 through 10 cover the arithmetic, the competitive consequence, the blockers and the execution.

Each chapter builds on the previous one. An executive who has their numbers buys what they need; one who doesn't buys what they are offered.

What you already pay for, and don't count

Before talking about tomorrow, today's cost has to be established. Not the technology budget — you know that one, it fits on one line. The full cost of moving information through your company, whatever department it is booked in.

Take a reference company: **120 employees, \$22M in annual revenue**, a plant or a warehouse, an office, an outside IT provider. You will adjust the numbers with your own; the structure does not change much from one sector to the next.

TABLE 1 — ANNUAL COST OF THE INFORMATION CHAIN, REFERENCE COMPANY

LINE	BASIS	ANNUAL COST	% REV.
S1 Customer support and after-sales service	4 FTE × \$60,000 loaded × 90% of time	\$216,000	1.0%
S1 Coordination, assistants, reception	4 FTE × \$69,000 loaded × 95% of time	\$262,200	1.2%
S1 Inside sales, estimating, quoting	5 FTE × \$88,000 loaded × 60% of time	\$264,000	1.2%
S1 Purchasing and procurement	3 FTE × \$78,000 loaded × 70% of time	\$163,800	0.7%
S1 Finance, invoicing, receivables	4 FTE × \$72,500 loaded × 80% of time	\$232,000	1.1%
S1 Quality, compliance, documentation	2 FTE × \$76,000 loaded × 70% of time	\$106,400	0.5%
S1 Human resources and payroll	2 FTE × \$75,000 loaded × 65% of time	\$97,500	0.4%
S4 Production planning	3 FTE × \$82,000 loaded × 75% of time	\$184,500	0.8%
S2 Outsourced managed IT	120 seats × \$115/month	\$165,600	0.8%
S2 Software licences and integrations	120 employees × \$1,900/yr	\$228,000	1.0%
S2 Website, digital marketing, media	Agency \$3,500/month + media buys	\$72,000	0.3%
S3 AI layer, all steps	A few individual subscriptions, marginal use	\$2,400	0.0%
S4 Uninvoiced internal coordination	12% of payroll of the other 93 employees (\$6,324,000)	\$758,880	3.4%
Visible total	All lines above except the last	\$1,994,400	9.1%
Real total	Internal coordination included	\$2,753,280	12.5%

Column S = step in the chain (see figure 1). Loaded cost = annual salary × 1.25. The share of time is the share spent moving information, not doing the trade itself: an estimator spends 60% of the week transcribing quantities, not estimating. Internal coordination is set at 12% of the payroll of employees *not* counted above, to avoid double counting – an assumption deliberately below published estimates (see appendix).

How many people are involved

The first finding of this exercise is the number of people involved. In a 120-employee company, **27 full-time equivalents** have moving information as their main function, spread across eight different roles, and none of them reports to the technology department. The estimator, the buyer, the payroll clerk, the planner: for most of their week, their job is to move information from one place to another, re-entering it along the way.

This company's technology budget, as it appears in the financial statements, is around \$466,000: managed IT, licences, the website and digital marketing. That budget is less than a fifth of the real chain.

The 27 FTEs are booked as salaries, spread across eight different cost centres. And the coordination of the other 93 employees — the time spent asking each other where a file stands — is counted nowhere.

This is the first brick to lay: **your real information cost sits mainly in your payroll, not in the IT budget.** As long as it stays scattered there, it cannot be optimized, because no manager is accountable for it.


BEFORE GOING FURTHER

Redo table 1 with your own numbers. Half a day with your controller is enough to produce a version defensible within 15%. That number — your real information cost as a percentage of revenue — is the only point of comparison that will make the next six chapters useful.

Four steps, one chain

Chapter 2 established what the chain costs. This chapter gives its structure. Information is acquired, stored and transmitted, processed, then converted into a decision: four steps, each with a real cost and its own unit of measure. As in any chain, the performance of the whole is limited by the weakest step, and column S in table 1 already ties each expense line to one of them.

FIGURE 1 — THE SEQUENCE: FROM RAW MATERIAL TO DECISION



Distribution computed from table 1, i.e. today. Steps 1 and 4, the ones that appear in no technology budget, weigh more than four fifths of the chain. The AI layer weighs one thousandth today because the mechanical work is still done by humans.

Step 1 — Acquisition: bringing information in

Everything you pay so that a piece of information from the outside world becomes a piece of information inside your company. The cost per lead in your digital channels. The order received by email that a clerk re-enters into the ERP. The PDF drawing an estimator reads by hand. The delivery slip photographed when the truck comes back. The supplier invoice entered a second time.

The useful units of measure are **entry cost per document** and **acquisition cost per customer served**, rather than the marketing budget. A distributor processing 900 orders a month with two order-entry clerks pays about \$12 per order just to transcribe what the customer already wrote. The customer pays for the product; the transcription has no value to them.

What the AI layer does here. Read the email, the PDF or the photo and extract the fields; sort what comes in and route it to the right place; filter incomplete requests and ask for what is missing. This is the step where human work is the most mechanical, so the first where the bill moves from salary to compute — and the first where a unit cost per document becomes measurable.

Step 2 — Storage and transmission: finding it when you need it

Managed IT, servers, cloud, software licences, integrations between systems, the website. This is the only step your company already names and budgets. It is also the one where the gap between what you pay and what you use is best documented: in large organizations, reserved cloud resources are massively underused, and taking back control of those costs has become a discipline in itself.

For a mid-sized company the practical question is this: **how many of your systems don't talk to each other?** Every boundary between two pieces of software that ignore each other is a place where a person bridges the gap. That cost is a people cost, booked in step 1 or step 4, not a licence cost.

What the AI layer does here. Find a file from a question asked in plain words rather than a reference number; link what is scattered across the ERP, emails and spreadsheets; file what was never filed. It does not replace your systems; it bridges them — which, until now, a person did by hand.

Step 3 — Processing: getting something out of it

The new step. Until recently, working on the content of the information itself — reading a drawing and pulling quantities off it, summarizing a call, drafting a first version of a contract, reviewing a quote, comparing three supplier offers — required a qualified human. The price of that processing was therefore the price of an employee-hour. Today there is a second price, measured in tokens, the billing unit of model providers.

The structure of that price is very different from a salary. An employee costs 40 hours a week whether they have too much work or too little; the bill is fixed and capped. AI is billed by usage. It never complains about having too much work: it does the work and bills "the overtime" — ten times, a hundred times the planned volume if demand calls for it, or if nobody set a limit. That is what makes this expense both profitable and one to watch closely: the volume consumed has no natural ceiling.

In table 1 the AI layer comes down to a few subscriptions, because the reference company has not yet handed AI anything. That is not its natural size. A company that genuinely hands it data entry, retrieval and decision preparation spends, at current public rates from model providers, \$5,000 to \$10,000 per month per equipped employee — several tens of thousands of dollars a month for a modest team. That spend replaces hours paid as salaries in steps 1, 2 and 4, so it pays off — provided its unit cost is managed with the same discipline as raw materials. Another risk is **putting a high-performance AI layer on top of a step 1 where the data is incomplete or wrong, and a step 4 where decisions sit waiting**. A company that automates processing without fixing acquisition produces fast answers from wrong data, and gets speed in the wrong direction. That is a net loss, just like a poorly trained employee who does not deliver the expected value — with one difference: billed by usage, this cost can run away much faster than a salary.

Step 4 — Coordination: turning it into a decision

The step that is never invoiced. Five people in a room to settle an equipment choice. Three follow-ups to get a number that already existed. The Monday morning production meeting. The approval waiting for the director to come back from vacation. The knowledge that leaves with the retiring partner.

What the AI layer does here. Prepare the decision file — the numbers, the options, the precedents — before people walk into the room; chase and track approvals without an assistant spending the day on it; retain what was decided and why. It does not decide. It

removes from the decision everything that is not the decision itself, and that is precisely where most of the cost sits.

Economist Ronald Coase put it this way in 1937: a firm exists because it is sometimes cheaper to coordinate internally than to go through the market. The corollary is rarely stated: **when your internal coordination cost rises, your company becomes objectively less justified in existing at that size.** That is the subject of the next chapter: this step is the most expensive of the four and the only one that cannot be bought.

The cost of a decision

You know what a machine-hour costs. You know what a truck-kilometre costs. The cost of a decision, on the other hand, is rarely known — even though the company produces hundreds of them a month and each one consumes the time of its best-paid people.

Management literature has sized the phenomenon from the top down: Bain estimates that the average company loses about a fifth of its productive capacity to what it calls organizational drag. McKinsey's work on decision-making lands on similar orders of magnitude. Those figures are useful at the scale of a large organization; they do not help run a mid-sized company. What is missing is **a simple method to price a given decision, made in your company.**

The proposed method fits on one page and requires no software.

DECISION COST CALCULATOR

$$\text{Cost} = (\Sigma \text{ participants} \times \text{hours} \times \text{loaded rate}) + \text{preparation} + \text{latency}$$

COMPONENT	HOW TO MEASURE IT	EXAMPLE
Meeting time	Number of meetings × duration × number of participants × each one's loaded hourly rate	\$2,040
Preparation and follow-up	Hours spent building the numbers, writing the minutes, chasing the absentees	\$1,190
Information retrieval	Hours spent finding what the company already knew but could not locate	\$680
Latency	Weeks elapsed between "we need to decide" and "it's decided", × the expected weekly gain from the decision	\$6,000
Cost of the decision	A single equipment purchase decision, six weeks, five people	\$9,910

LOADED HOURLY RATES, TO ADJUST

Annual salary × 1.25 ÷ 1,800 h · Owner \$180,000 → \$125/h · Director \$110,000 → \$76/h · Professional \$85,000 → \$59/h · Coordinator \$60,000 → \$42/h

Ten thousand dollars for one decision. Applied to your company's last five structuring decisions, the exercise produces an amount that finally puts step 4 on the agenda of a management meeting.

Note what the formula says and does not say. It does not conclude that meetings are useless, or that one should decide alone and fast. It isolates **the "information retrieval" and "latency" components**, which do nothing to improve the quality of the decision. That is where AI comes in: it does not decide, it lowers the cost of finding what the company already knows.

And that is where the loop closes back on Coase. The higher your information cost per decision, the more vulnerable you are to a smaller competitor whose cost is close to zero. That competitor does not need to be better at manufacturing or service; they only need to decide ten times while you decide once.

Add to that the case where this competitor uses agentic AI to set up continuous improvement loops on their chain: agents that watch the outcome of each decision and adjust the next one. Favouring the most profitable acquisition channel over another. Changing the keywords bought on search engines based on what actually converts. Adding to the quote form the two precise questions that were missing and forced a callback to the customer. Each of these adjustments is small; none requires a meeting. The competitor thus becomes more efficient than you on a regular basis, week after week, and **the gap ends up showing in a pricing dynamic that becomes very hard to catch up to.**

Mapping your chain like an org chart

You have an HR org chart. It shows redundancies, gaps and lines of accountability. Nobody reorganizes a department without looking at it.

There is no equivalent for information. That is why a mid-sized company can have three people producing the same report, an inventory system that customer service never checks, and nobody able to answer "what does a quote cost us?". The work to do is the same: **draw the org chart of the cost of information.**

TABLE 2 – MAPPING GRID, ONE ROW PER FLOW

COLUMN	THE QUESTION IT ANSWERS	EXAMPLE: THE "QUOTE" FLOW
Flow	What information comes in, and what must it become?	Request for pricing → signed quote
Entry point	Where does it arrive, and in what form?	Email, PDF drawings, phone call
Hands	How many different people touch it?	5: sales, estimating, production, purchasing, finance
Time	How many human hours per unit processed?	4.5 h per quote
Unit cost	Hours × loaded rate + attributable licences	\$340 per quote
Latency	How long between entry and decision?	4 business days
Owner	Who, by name, is accountable for this flow's cost?	Director of operations

One row per information flow, not per piece of software. Eight to fifteen flows are enough to cover a 50-to-200-employee company.

How to fill in the grid

Start from flows, not software. A flow is a sentence like "a request for pricing comes in and becomes a signed quote". It crosses sales, estimating, production and finance, and that is exactly why it is expensive. A software inventory helps renegotiate licences; it does not recover margin.

Count the number of people who touch the information. For each flow, the most revealing question is: how many different people touch this information before a decision is made? Beyond four, the problem almost always sits in step 4.

Name a cost owner for each flow. The person who can answer "a quote costs us \$340" and who can be asked to bring that number down. Without an owner, the map stays a document.

Spot the places where a person bridges two systems. Every time someone takes information from one system to re-enter it into another, that is a cost with no value to the customer. Those points are the first to address: they are measurable, isolable, and fixing them requires no delicate reorganization.

Six sectors, broken down

The six tables that follow describe the same transformation in six contexts. They read left to right: the function, what it costs today in a typical mid-sized company, what it looks like once the information chain is rebuilt, and where the gain shows up in your financial statements.

In none of these tables does the "tomorrow" column describe replacing a person. It describes a person who stops moving information and goes back to their trade. An estimator entering quantities is transcribing, not estimating.

6.1 Manufacturing

MAKE-TO-ORDER, 80 TO 200 EMPLOYEES

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Estimating	An estimator works through PDF drawings by hand, 3 to 6 h per quote; a good share will never be won	Quantities and bill of materials extracted from drawings; the estimator validates and sets the price, 45 min	Revenue ↑
Planning	Shared spreadsheet, Monday production meeting, 6 people, 90 min	Production schedule recomputed continuously; the meeting settles exceptions, 25 min	Coordination ↓
Quality	Non-conformance reports on paper, reviewed quarterly	Each non-conformance tied to machine, lot and shift, continuously	Losses ↓
Purchasing	Three emails and two calls per part; price negotiated from memory	Consolidated price and lead-time history, purchase recommendation prepared	Material cost ↓
After-sales service	A technician searches the files to find the configuration that was shipped	Full history of the delivered equipment, available in seconds	Billable hours ↑

6.2 Distribution and wholesale

WAREHOUSE, CATALOGUE OF SEVERAL THOUSAND PRODUCTS

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Order entry	Orders by email and phone, re-keyed by hand into the ERP	Automated reading and entry; the clerk handles exceptions	FTEs reassigned
Pricing and margins	Fixed price lists, discounts granted at the rep's discretion	Margin per customer and per line visible at quoting time	Margin points ↑
Inventory	Cycle counts; stockouts discovered by the customer	Forecast per product and season, alert before stockout	Lost sales ↓
Customer service	Most calls are about order status	Order status available without a call; people keep the disputes	Cost per contact ↓
Receivables	Manual follow-ups when someone remembers	Sequenced follow-ups prioritized by risk	Working capital ↑

6.3 Professional services

ACCOUNTANTS, LAWYERS, ENGINEERS, 50 TO 150 PROFESSIONALS

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Producing the document	Drafted from an old file someone recalls from memory	First version built on the firm's precedents; the professional reviews and signs	Non-billable hours ↓
Billing	Timesheets filled in on Friday, from memory, under-reported	Time reconstructed from work traces, validated by the professional	Recovered revenue
Intake	Web form, callback within 48 h	Qualification and documents requested at first contact	Conversion ↑
Research	A junior spends a day finding the applicable case law or standard	Research prepared and sourced, reviewed in 45 min	Margin per mandate ↑
Internal knowledge	Knowledge leaves with the retiring partner	Files, opinions and precedents searchable	Risk ↓

6.4 Construction

GENERAL OR SPECIALTY CONTRACTOR, SEVERAL CONCURRENT SITES

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Bids	The estimator builds four bids a month; two are won	Quantities and prices built from drawings; bid volume triples	Backlog ↑
Site tracking	Handwritten daily report, entered the following week	Report dictated on site, structured and tied to the lot the same day	Billed change orders ↑
Extras	Discovered at final invoicing, then disputed	Variance detected and documented the day it happens	Recovery ↑
Labour	The foreman spends two hours a day on the phone reassigning people	Assignment proposed; the foreman decides	Coordination ↓
Compliance	Binders, searched on demand when the inspector shows up	Compliance file kept up continuously	Penalties ↓

6.5 Transportation and logistics

FLEET OF 20 TO 100 VEHICLES, GENERAL FREIGHT

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Dispatch	The dispatcher plans by eye, over the phone	Routes optimized continuously; the dispatcher handles the unexpected	Empty km ↓
Documents	Paper delivery slips, entered when the truck returns	Photo of the delivery slip read and tied to the trip	Invoicing delay ↓
Pricing	Per-kilometre grid, route prices reviewed once a year	Real cost per trip, per route and per customer	Money-losing trips ↓
Customer service	"Where is my truck?", by phone, all day	Trip status available without a call	Cost to serve ↓
Maintenance	Inspections recorded on paper, reactive maintenance	History per vehicle, maintenance planned on real signals	Downtime ↓

6.6 Food processing

PROCESSING, RETAIL AND DISTRIBUTOR CUSTOMERS

FUNCTION	TODAY	TOMORROW	WHERE THE GAIN GOES
Traceability	Lots tracked in a spreadsheet; a recall takes three days of searching	Upstream-downstream traceability searchable; targeted recall in hours	Risk ↓
Yield	Losses found in the monthly statement, no cause identified	Material variance per shift and per recipe, continuously	Losses ↓
Retailer orders	One customer portal per chain, filled in by hand	Automated filling and reconciliation	Penalties ↓
Labelling	Spec sheets and allergens redone at every recipe change	Sheets regenerated from the source recipe	Errors ↓
Forecasting	Last year's history and intuition	Forecast per product, promotion and season	Overproduction ↓

WHAT THE SIX TABLES HAVE IN COMMON

Across the six sectors, the same three patterns recur. **One:** the most profitable function to address first is the one that touches price — estimating, quoting, pricing, margin per line. That is where a few hours of recovered human work convert directly into points of margin. **Two:** the second most profitable is answering "where does my file stand", which accounts for the majority of inbound contacts in almost every service business. **Three:** compliance and traceability do not cut costs, they cut risk — and that is often the argument that gets the project past a cautious board.

These three patterns are cost projects, led by executives or finance, not IT projects.

The only two equations that matter

All the complexity of the subject comes down to two paths. The company chooses which one it takes, department by department, before buying anything.

EQUATION A – THE COST LEVER

Revenue constant
Costs go down
→ Profit goes up

The usual path, and the easiest to defend to the board. You keep the same volume of business and remove from the chain the work the customer does not pay for: entry, re-entry, information retrieval, follow-ups.

The trap: a company that takes only this path becomes more profitable without becoming stronger. Its costs fall; its competitive position does not change. A year later, the competitor who took path B takes its customers.

EQUATION B – THE CAPACITY LEVER

Costs constant
Revenue goes up
→ Profit goes up

The more demanding path, and the more profitable. Nobody leaves. The freed capacity is reinvested: the estimator produces three times as many quotes, the rep covers more territory, the firm accepts mandates it used to turn down.

The condition: there has to be available demand. In a saturated market, or one constrained by production capacity, lever B does not work, and that has to be acknowledged before promising it to the teams.

One lever per department

The most common mistake is choosing one path for the whole organization. The two do not apply in the same places. Finance and administration almost always fall under lever A: the volume of accounting entries does not need to triple. Sales, estimating and service almost always fall under lever B: capacity freed there converts into revenue, provided there is a market.

TABLE 3 – WHICH LEVER, IN WHICH DEPARTMENT

DEPARTMENT	LEVER	WHAT CHANGES CONCRETELY	METRIC TO TRACK
Sales and estimating	B	The same headcount produces two to three times as many quotes, and answers faster	Quotes/month, response time
Customer service	A then B	Repetitive requests no longer consume human time; the team moves to proactive service	Cost per contact
Finance and administration	A	End of invoice re-entry and manual reconciliation	Cost per invoice processed
Operations and production	A	Fewer coordination meetings, less time spent looking for a file's status	Meeting hours/week
Purchasing	A	Consolidated price and lead-time history, fact-based negotiation	Material cost per unit
Human resources	A	Files, training and compliance kept without manual handling	Admin hours/employee
Executives	B	The cost of knowing what the company knows falls; the pace of decisions rises	Cost and latency per decision

A department without a numeric metric falls back into its old habits within six months.

Now add it up. If your real information cost is twelve percent of revenue — the order of magnitude in table 1 — and you remove a third of it without touching business volume, you have just recovered four points of net margin. For most 50-to-200-employee companies, four points of net margin is more than all the productivity gains of the past ten years combined.

Why the gap doesn't close

A prudent executive thinks: "I'll look at this in two years, once it's stable, and I'll catch up in six months." That reasoning served well for previous technologies. The website, the ERP, the cloud: in each case, waiting cost little and the latecomer bought the same thing as the pioneer, often for less.

Here the reasoning does not hold, for a structural reason: **what accumulates is the organization around the tool, not the tool itself.** The competitor who starts now is building a clean information base, flows that someone is accountable for, teams that have learned to work with a low decision cost. Those three things take calendar time and cannot be bought. This is what we call the agentic AI transformation.

THE MECHANISM, IN FOUR MOVES

1. They lower their cost to serve

Their quotes cost \$90 instead of \$340. Their cost per customer contact falls by half.

2. They choose what to do with the gain

They can keep it as margin, or give it back to the customer in price and lead time. In a contested market, they give it back.

3. You lose bids

On price and on response time. You conclude it is a price war and cut spending elsewhere.

4. The gap widens

Their volume gives them more data, hence an even lower cost to serve. You cut what would have let you catch up.

The third move is the most damaging, because it is misdiagnosed. An executive who loses three bids in a row on price concludes the market has deteriorated. They cut the sales budget, postpone investments, squeeze suppliers. Those are the right responses to a price war, and the wrong ones against a competitor with a better cost structure.

The signal to watch is **competitors' response time**, before their price. A company that answers a request for pricing in four hours where you take four days has, in most cases, rebuilt its information chain rather than hired estimators. Lead time is the first symptom visible from outside, and it appears months before price moves.

THE ADVANTAGE OF SIZE

A 50-to-200-employee company is *advantaged* on this ground. You do not have forty legacy systems to reconcile, a governance committee to convince, or thousands of positions to requalify. You have a management meeting where everyone fits in the same room and an owner who can decide on Tuesday morning. That is a real structural advantage over the large companies in your sector, and the window to use it is short.

The most common blockers

These three come up in almost every conversation with owners of mid-sized companies. All three are legitimate. Each changes the order in which the work should be done.

"My data isn't ready. Everything is in spreadsheets and in people's heads."

That is accurate. The blocker, however, rests on a premise that no longer holds: that you need a clean database before starting. What was true in the era of data warehouse projects is no longer true. A poorly kept spreadsheet, an email, a photographed handwritten note are now processable inputs. The constraint has shifted from structure to accuracy.

The practical answer: **pick a flow where the information is already reliable, even if it is poorly structured.** Your supplier invoices are accurate. Your delivery slips are accurate. Your drawings are accurate. Those flows are ready today. The flow "who told the customer what last week" is not, and it can wait.

And on "in people's heads": that is precisely the risk the mapping in chapter 5 makes visible. A company where three people each hold undocumented knowledge has a business continuity problem first. The work to do is the same.

"I have nobody in-house to lead this."

The role required is not technical. The first ninety days described in the next chapter require no IT skills. They require someone able to keep a cost spreadsheet, get people from different departments talking, and say no to projects that cannot be measured. In most mid-sized companies, that person exists: the controller, the director of operations, or the owner.

Handing the file to the IT provider for lack of a better option is, on the other hand, a common mistake. A managed IT provider is competent on step 2 — that is their trade, and they do it well. They have no mandate, no tool and no interest in measuring your internal coordination cost. Asking them to run the whole chain is like asking your carrier to organize your production.

The rule of division is simple: **steps 1 and 4 stay with you, because they are your business. Step 2 and the AI layer can be bought.** You do not run your own data centre; you do not need to run your own AI models either. Nobody outside, however, can decide which flows matter and who is accountable for them.

"My sector is too specialized, or too regulated."

Specialization is a serious argument when it concerns *judgment*. An engineer who signs a drawing, a veterinarian who makes a diagnosis, an estimator who senses a customer is about to change their mind: that judgment is not reproducible and is not at issue here. Look at the six tables in chapter 6: the "tomorrow" column never replaces the trade's judgment; it removes the hours spent preparing the conditions for it.

As for regulation, it often works the other way. A regulated sector is one where traceability, documentation and reporting cost a great deal in human hours — so one where steps 1 and 4 are particularly inflated. Food processing and construction, in tables 6.6 and 6.4, are the best examples: compliance there is both the heaviest line and the most mechanical.

What regulation actually imposes are constraints on **step 2 and the AI layer**: where information and data reside, who accesses it, what is logged, what may leave the country. Those constraints are technical and are settled when choosing a provider.

A fourth blocker often comes up: "I tried, my employees don't use it." A tool handed to a team without any flow being redesigned goes unused, and that is a rational response. Behaviour changes when valueless work is removed, not when licences are distributed.

Ninety days, then twelve months

The first quarter deliberately contains no purchase. At the end of ninety days, the company has its own numbers and can evaluate a vendor proposal on that basis.

TABLE 4 – THE FIRST NINETY DAYS

PERIOD	WHAT YOU DO	EXPECTED RESULT
Days 1-15	Redo table 1 with your numbers, with your controller. No software, no vendor.	Your real information cost, as % of revenue
Days 15-30	Map eight to fifteen flows with the table 2 grid. One meeting per department, 90 min.	Org chart of the cost of information
Days 30-45	Price three recent decisions with the chapter 4 calculator. Pick real decisions, not examples.	Average cost of a decision in your company
Days 45-75	Pick a single flow: the one whose unit cost is best measured. Name its cost owner.	One flow, one owner, one starting cost
Days 75-90	Redo that flow with agentic AI: the agent enters, retrieves and prepares; the person validates and decides. Measure afterwards, on the same basis as before.	A before/after gap defensible to the board

Actual load: about one person-day per week, plus 90 minutes per department in week 3.

TABLE 5 – THE TWELVE MONTHS, WITH EXIT CRITERIA

QTR	OBJECTIVE	WHEN TO STOP
Q1	The first flow is in production and its gain is measured. Table 1 becomes a quarterly tracker, not an exercise.	If the gain is not measurable: stop and change flows
Q2	Two more flows, chosen in two different departments to test the method elsewhere.	If the second flow takes longer than the first: the method is at fault
Q3	Explicitly choose lever A or B for each department, with a numeric metric per department.	If a department has no metric: it is not ready
Q4	Reinvest the freed capacity. Revisit step 1 and 2 contracts in light of the accumulated numbers.	If information cost as % of revenue has not fallen: redo the diagnosis

An information chain project that produces no measured gap in one quarter will not produce one in four.

Three common causes of failure

Starting with the biggest flow. The most expensive flow is almost always the most delicate to touch. Start with the one whose unit cost is easiest to measure, even if it is small: the goal is an indisputable before/after gap.

Running a pilot without a cost owner. A pilot led by "the IT team and a consultant" produces a demonstration. A pilot led by the person accountable for the flow's cost produces a measured gap. The difference cannot be recovered at the end of the project.

Not deciding what to do with the freed capacity. If you take eight hours a week away from an estimator without telling them what to do with those eight hours, they fill up on their own within three weeks and you will have no gain to show. Lever A or lever B is chosen *beforehand*.

Where Agent Libre fits

This document has an interested author; that should be stated.

Agent Libre offers a managed inference service. In other words: **we provide the AI layer — the compute it consumes across all four steps — and nothing else.** We run the models and the infrastructure behind them on behalf of mid-sized companies that do not want to manage compute servers, negotiate reserved volumes, or watch a unit price that changes every quarter. You hand us the compute; you keep control of your flows, your data and your decisions.

WHAT WE DO

- Run the AI layer: compute capacity, availability, cost per unit processed, processing history
- Give you a predictable, verifiable AI cost, line by line, whichever step consumes it
- Track on your behalf the changes in models, prices and providers that happen every quarter, without your bill or your service being affected

WHAT WE DON'T DO

- Map your organization for you — steps 1 and 4 are your business
 - Replace your managed IT provider on step 2
 - Provide the specialized software your employees use and that calls the language model
 - Decide which flows matter, or who is accountable for them
-

This boundary follows from the document's argument. A vendor offering all four steps creates a dependency. Step 4 — the cost of your decisions — is the only one that belongs entirely to the company, and it should not be outsourced.

The test we propose to our clients is therefore the same as in chapter 10: do the ninety days without us. If, at the end, usage-based AI billing is not a problem, our services are not needed. If it is, you will know what you are buying and at what cost per unit.

WHAT TO TAKE AWAY

The most important question is not the price of subscriptions or the usage bill, although that matters too. It is "what does moving information through my company cost me today, and what share of that cost is the customer willing to pay for". For this document's reference company, the answer is \$2,753,280 a year, or 12.5% of revenue, of which more than four fifths appear in no technology budget. It is the single most important line an owner of a mid-sized company has to optimize over the next decade, and optimizing it is built rather than bought.

Sources and methodological notes

ON COORDINATION AND DECISION COST (STEP 4)

Ronald H. Coase, "The Nature of the Firm", *Economica*, 1937. Theoretical foundation of internal transaction costs. Qualitative; no per-decision figure.

Oliver Williamson, work on transaction costs, 1975 and 1985. Extends Coase on the boundaries of the firm.

Michael Mankins and Eric Garton, *Time, Talent, Energy*, Bain & Company / Harvard Business Review Press, 2017. Source of the notion of organizational drag, estimated at about a fifth of the average company's productive capacity, based on a survey of more than 300 executives conducted with the Economist Intelligence Unit. Macro-level quantification, not directly applicable to a single decision.

Aaron De Smet, Gregor Jost and Leigh Weiss, "Decision making in the age of urgency", McKinsey & Company, 2019. Survey of more than 1,200 executives.

ON PROCESSING AND INFRASTRUCTURE COST (STEPS 2 AND 3)

ICONIQ Growth, *State of AI*, 2026. Primary data from about 300 technology company executives: gross margins of AI products, share of processing in the cost structure. Public, free.

Bessemer Venture Partners, "The AI pricing playbook for founders", 2026. Documents the gross margin gap between AI products and traditional software. To be attributed institutionally.

Sarah Wang and Martin Casado, "The Cost of Cloud, a Trillion Dollar Paradox", Andreessen Horowitz, 2021. Origin of the argument that every point of gross

FRENCH-LANGUAGE AND QUEBEC CONTEXT

Cigref, "Repenser la valeur des projets de transformation numérique pour l'entreprise", 2018. Methodology for measuring the value and hidden costs of digital projects. Oriented toward project value, not the cost chain.

CAVEATS

The extrapolations in days and dollars are presented by the authors themselves as a thought experiment, not a measurement.

Steven Rogelberg (UNC Charlotte) and Otter.ai, "The Cost of Unnecessary Meeting Attendance", 2022.

Survey of 632 employees across 20 sectors. Reproducible method (loaded hourly rate × hours) but at the meeting level, not the decision level. Vendor-sponsored study.

Klein and Wieczorek, "The Headless Firm", *arXiv:2602.21401*, 2026. Models the fall in coordination cost and its effects on the boundaries of the firm. Theoretical, with no operational management dimension.

margin recovered on infrastructure flows through to company value.

FinOps Foundation, *State of FinOps 2026*. Survey of 1,192 practitioners. Establishes that managing AI cost has become near-universal among organizations that manage cloud spend. No organizational layer in its framework.

Cast AI, *State of Kubernetes Optimization*, 2026.

Very low utilization rates of reserved resources, including GPUs. Proprietary data from a vendor that sells optimization: commercial bias to note.

Statistics Canada and Institut de la statistique du Québec. Data on business adoption of AI and reported barriers, where high cost is cited as the first obstacle. No unified cost framework.

The amounts in table 1, the chapter 4 calculator and the sector tables are illustrative constructions, built on the reference company and on explicit loaded hourly rates. They are neither a market benchmark nor a commitment to results.

The assumption of 12% of payroll attributed to internal coordination is deliberately below published estimates, so that the demonstration does not depend on the most favourable figure.

Statements about the future trajectory of processing cost are projections drawn from observed trends, not established facts. A reversal of that trend would change the priority order recommended in chapter 10, without invalidating the four-step framework.

At the time of publication, no published framework — consulting firm, vendor, foundation or academic — brings the four steps together in a single model aimed at executives. Should one appear, it should be preferred to this one wherever it is better documented.