

THE OPERATING PLAYBOOK

UPDATED JUNE 2026

The 90-Day Protocol

The first ninety days that separate the 5% from the 95%. The report named the disease. This is the field guide for the cure — run on a calendar, in order, every quarter.

5% / 95%

The base rate is settled. The method is the part that moves you across the line. Source: MIT NANDA, 2025.

Diagnose. Fix the base. Compound.

A guide you keep open on the desk. The sequence is the whole argument.

BLANKCOLLAR.UNIVERSITY

THE EXECUTION WING

A protocol, not a pilot. A move you can start on Monday.

Three phases, six moves, ninety days. The figures are the report's figures — MIT, McKinsey, IBM, Cursor, and the three autopsies — kept straight and put to work. Read it once, then run it.

01	Why a protocol beats a pilot The learning gap, the 95%, and the cost of work about work	p . 03
02	Phase one — Diagnose · Days 1–30 Score the Stack, name the red layer, cancel the theatre	p . 05
03	Phase two — Fix the base · Days 31–60 Raise the base above one before you turn on the roof	p . 06
04	Phase three — Compound · Days 61–90 Redeploy, hire the orchestrator, re-score, repeat	p . 07
05	The six moves, expanded Buy, Aim, Clean, Keep, Redeploy, Cancel — with worked examples	p . 08
06	The Stack as the operating model Shrink the routine process, grow the human experience	p . 12
07	The three traps Speed without direction, literacy as paperwork, ROI theatre	p . 13
08	What to measure — and what not to claim Hours, processes, rungs — never a universal ROI	p . 14
09	The 90-day calendar Twelve weeks, week by week, fillable	p . 15
10	Direct, don't do The close, the two moves, and the sources	p . 16

A pilot asks whether AI works. A protocol decides whether yours lands in the 5%.

The report ended with a finding and a sequence. In 2025, MIT's NANDA team measured the profit-and-loss impact of corporate AI transformations at scale, and ninety-five per cent of them returned nothing a finance team could find. The failure did not trace to weak models or a shortage of talent. It traced to a **learning gap** – the tools could never absorb how the company actually works. This guide is the part that comes after the diagnosis: the operating playbook for the ninety days that move you off the base rate.

A pilot is a question asked in public. It runs for a quarter, produces a deck, and proves only that the model can do the thing in a sandbox – which was never in doubt. A protocol is different. It does not ask whether AI works; it fixes the base the AI multiplies, in a fixed order, on a fixed calendar. The distinction is the whole report in one line: **the model was never the variable. The layer beneath it was.**

The reason the order matters is arithmetic. AI is not an ingredient you add; it is an exponent, and an exponent on a base below one makes small things smaller, faster. Run the moves out of order – turn on the roof before the foundation holds – and you industrialise the mess at speed. Run them in order and the same exponent compounds. Ninety days is long enough to fix one layer and prove it moved; short enough that nobody can hide.

THE PREMISE

Most AI work should be cancelled, not improved – until the base it multiplies is worth multiplying.

95%

of corporate AI transformations returned zero measurable P&L impact. MIT NANDA, 2025.

58%

of the knowledge-worker day goes to work about work – the tax the protocol is built to cut. Cited in *The Blank Collar*.


AI does not subtract. It multiplies whatever you point it at.

THE BLANK COLLAR · 2ND EDITION

Everyone is in the gym. Almost no one is getting stronger.

The failure does not look like failure. It looks like activity. Two numbers that should move together have come apart — usage is near-universal, results are rare — and that gap is exactly what a protocol is for.

THE TWO CURVES THAT CAME APART



Adoption and earnings impact: McKinsey, 2025. P&L win: MIT NANDA, 2025.

95%

RETURNED ZERO P&L
IMPACT

5%

A MEASURABLE,
REPEATABLE RETURN

READ THE GAP, NOT THE DASHBOARD

If adoption alone worked, 78% would be winning. The result lives in the method, not the tool.

A 95% failure rate is not a performance problem you fix with another workstream. It is a mortality rate. The five who survived it did not get lucky and did not buy a better model. They followed an order you can read off them — and the rest of this guide is that order, set to a ninety-day clock.

DAYS 1-30

Diagnose

PHASE 01 / 03

No competent doctor prescribes before they diagnose. Month one buys you the truth: where you actually stand, which layer is red, and which initiatives are theatre dressed as progress.

THE GOAL

End the month knowing your red layer and your revenue per employee – the two facts that decide the next sixty days.

Score the five-layer Stack honestly: true on a bad week, not on board day. Read the shape, not the total. A company scoring 38 on AI and 12 on Data does not have a strong AI capability. It has a fast hallucination machine, and its real score is the 12.

THE EXIT TEST

You can name one red layer, one number a CFO can find, and the date you will move it.

If you cannot name the layer or the number, you are not ready to fix anything yet. Stay in diagnosis another week. A method that skips this step is a pilot wearing a protocol's clothes.

THE MOVES THIS MONTH

- Score the Stack.** Grade Vision, Data, Process, Human Experience and AI. Find the lowest. It caps everything above it.
- Compute revenue per employee.** Write it beside the benchmarks: median software ~\$130K, AI-native firms \$3-7M.
- Aim at one P&L line.** Name the single number, with a date, that this quarter will move.
- Cancel the theatre.** Any initiative with no P&L line is finished. Sign the certificate; free the budget.
- Do the diagnosis in-house.** No vendor has an incentive to find your dull, real problem. Walk in already knowing it.

WORKED EXAMPLE

A 400-person services firm scores **34 on AI, 14 on Data**. The instinct is to buy a better model. The protocol says the score is the **14**: stop funding the roof, and point month two at the data layer.

Most AI transformations should be cancelled, not improved.

THE THESIS

DAYS 31-60

Fix the base

PHASE 02 / 03

Month two raises the base above one. You do not turn on the multiplier yet – you make the thing it will multiply worth multiplying. One layer, the red one, fixed enough to prove it moved.

THE GOAL

Run the red layer's one 30-day move, and stand up a single small agent pilot from where you sit – inside policy.

Encode one process that lives only in a head. Free your three data hostages. The act of writing the process down is half the value: you cannot automate a mess, a vibe, or a habit with good PR – only what you have made explicit.

THE EXIT TEST

One process is written and firing. One number that was contested now has a single source. One agent ran two weeks against an eval.

If nothing fired automatically, you documented but did not wire it. A binder is not a reflex. Wire one handoff before the phase closes.

THE MOVES THIS MONTH

- Run the red layer's move.** Whatever month one named – the one 30-day fix for the weakest layer.
- Clean, then encode one process.** Excavate input, action, output. Run the five-whys. Delete the ghost steps.
- Free three data hostages.** The three facts only one person holds. Move each into the system, not a head.
- Run one agent pilot, in policy.** Minimum tools, a human approving anything consequential, a small eval.
- Keep a human where the brand lives.** Put a person on the moment that carries the brand, before you scale it.

WORKED EXAMPLE

The services firm encodes its month-end close – a **57-step approval** that the five-whys reveals to be three real steps wearing fifty-four ghosts. The agent now reconciles the routine; a person signs the **exceptions**.

You cannot automate a mess. You cannot automate a vibe. You can only automate what you have written down.

ON THE PROCESS LAYER

DAYS 61-90

Compound

PHASE 03 / 03

Now you turn on the multiplier. Month three redeploys what you saved, hires for the new shape of work, and re-scores the red layer. If it moved, you have proof – and a method you can run again.

THE GOAL

Redeploy the reclaimed hours into work that compounds, hire one orchestrator not one operator, and re-score the layer you fixed.

The hours you freed are the point of the exercise. IBM automated roughly 94% of routine HR tasks and total employment rose – because the hours moved up the value chain, not off the payroll. Point the freed time at judgment, novelty and trust.

THE EXIT TEST

The red layer's score moved. You can show one P&L line that changed, and name the next quarter's red layer.

If the score did not move, the protocol still worked: it told you the truth cheaply. Re-aim and run it again. A method you can repeat is the asset, not any single win.

THE MOVES THIS MONTH

- Redeploy the savings.** Move the reclaimed hours to judgment and novelty – not to a smaller payroll.
- Hire the orchestrator.** The next role designs and runs the system. It does not perform the task by hand.
- Re-score the red layer.** Run the same audit. Did the number move? That is your proof, or your re-aim.
- Read the P&L line.** Show the single number from month one. A CFO should be able to find it.
- Set the next ninety days.** Name the new red layer. The protocol is a loop you run every quarter, forever.

WORKED EXAMPLE

With the close automated, the firm redeploys two analysts onto client advisory and re-scores Data from **14 to 22**. The next red layer is now Process – and the **next ninety days** have a target.

Diagnose. Fix the base. Compound. Then run it again – the loop is the asset.

THE 90-DAY PROTOCOL

The six moves of the 5%, expanded.

None of these needs a bigger model. Every one is a decision available today. The next four spreads take them one at a time — what it is, when it lands in the ninety days, and a worked example.

01 Buy before you build.

WHAT IT IS Buy the capability that already works. Bought tools reach production about **twice as often** as in-house science projects. Do not fund a research department by accident.

WHEN Phase one and two, the moment a need is named. The default is buy; building is the exception you must justify.

WORKED EXAMPLE

A team scopes a six-month internal copilot. The protocol kills it: a bought tool covers **90%** of the need in a fortnight, and the saved months go to the data fix that actually gates the win.

02 Aim at one number.

WHAT IT IS Aim at a single P&L line a CFO can find — not “a transformation.” **One number, one date.** If you cannot name it, you have a budget line for theatre.

WHEN Phase one, before a cent is committed. The aim is the contract the whole ninety days is judged against.

WORKED EXAMPLE

“Improve support with AI” becomes **cost-per-resolved-ticket down 30% by day 90, satisfaction held flat.** Now the pilot has a **pass/fail**, not a press release.

03 Clean the process.

WHAT IT IS Clean the process before you automate it. Write it down, run the **five-whys**, delete the ghost steps. You can only automate what you have made explicit.

WHEN Phase two. Cleaning is the work that raises the base above one. Automate a mess and the exponent industrialises it.

WORKED EXAMPLE

A **57-step approval** turns out to be three necessary steps wearing fifty-four ghosts. Delete the ghosts first; the survivor is trivial to wire and an agent runs it without supervision.

04 Keep a human.

WHAT IT IS Keep a human where the brand lives. Accountability cannot be delegated to a tool. Put a person on the moment that carries the brand and owns the verdict.

WHEN Phase two and three, wherever an agent touches a customer. The human moves up the loop; it does not leave it.

WORKED EXAMPLE

Air Canada's chatbot invented a refund policy and a tribunal held the airline to it. The fix is not less AI — it is a named human who **owns what the agent says** at the edge.

NOT

Not a sign-off committee. Keeping a human means one accountable person on the brand-bearing moment — not a queue of approvers that turns the agent back into a meeting. Reliability is a discipline, not a head-count.

05 Redeploy the savings.

WHAT IT IS Redeploy the savings into work that compounds. The freed hours go up the value chain – to judgment, novelty and trust – **not** into a smaller payroll.

WHEN Phase three. The point of the saving is what you point it at next; aim it at a headcount cut and you get a smaller company.

WORKED EXAMPLE

IBM automated **~94% of routine HR tasks** and total employment **rose** – the hours were redeployed, not deleted. Aimed well, AI does not subtract; it **multiplies**.

06 Cancel fast.

WHAT IT IS Cancel fast. A dead initiative is hygiene, not shame. Any pilot that cannot survive moves one to five is already finished – **sign the certificate** and free the budget.

WHEN Every phase. Cancellation is the move that funds the others; a hope without a deadline is a way to postpone the decision while the best people leave.

WORKED EXAMPLE

Klarna replaced human support to chase the cheapest answer, declared victory, then quietly rehired as quality fell. The cheapest answer was the most expensive one. **Cancelling early** would have cost a quarter, not a brand.

A dead initiative is hygiene, not shame. Sign the certificate, and free the budget for the one that can survive.

ON MOVE SIX

Six decisions, in order. None of them needs a bigger model.

01 Buy before you build.

Bought tools reach production about twice as often. Buy the capability that works; do not fund a research department by accident.

02 Aim at one P&L number a CFO can find.

Not a transformation. A single line, with a date. No number means no project — only a budget line for theatre.

03 Clean the process before you automate it.

Write it down; run the five-whys; delete the ghost steps. You can only automate what you have made explicit.

04 Keep a human where the brand lives.

Accountability cannot be delegated to a tool. Put a person on the moment that carries the brand.

05 Redeploy the savings into work that compounds.

Point the freed time at judgment, novelty and trust — not at a smaller payroll. IBM's headcount rose.

06 Cancel fast.

A dead initiative is hygiene, not shame. Anything that cannot survive moves one to five is already finished.

Nobody is born into the five percent. You choose your way in.

THE BLANK COLLAR

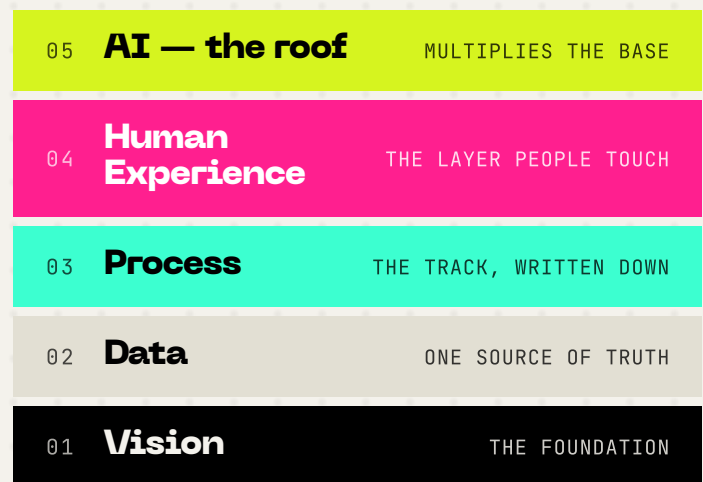
The Stack is the model the protocol runs on. You break at the layer you ignored.

The protocol is not a checklist floating free – it runs on one drawing. Five layers, bottom to top, where one thing must be true before the next can be. **Vision** is the foundation: one sentence sharp enough that a person or an agent can reject a bad project without asking. **Data and Process** are the load-bearing middle. **Human Experience** is the layer people touch. And **AI is the roof** – not a fifth box, but the layer that multiplies whatever is underneath.

That placement is the diagnosis in one stroke of a pen. A perfect roof on no foundation is 0% there, not 60%. The lowest layer caps everything above it, which is why month one finds the red layer and month two fixes it – before month three ever turns on the multiplier.


The one move the whole school repeats sits inside this drawing: **shrink the routine process, grow the human experience**. Hand the repeatable 99% to the machine; expand the 1% only a person can do. The protocol is that move, run on a calendar.

BOTTOM-UP: STABILISE BEFORE YOU TURN ON THE ROOF



THE ONE MOVE, DRAWN

PROCESS — SHRINK IT ↓



The routine 99% – drafting, reconciling, classifying, routing. Hand it to the machine; it gets cheaper as machines do.

HUMAN EXPERIENCE — GROW IT ↑



The final 1% – judgment, the relationship, the decision about what is worth building. It expands as the machine takes the rest.

Three traps sink more protocols than any model ever did.

Each one feels like progress from the inside. Each one is the report's own failure mode wearing a respectable face. Name them now, and you can see them coming in week three instead of month six.

01 **Speed without direction** Symptom: the team ships faster every week, but no one can name the P&L line it moves. Motion is mistaken for leverage.

THE CORRECTION

A faster mess is still a mess. Aim at one number before you optimise for velocity. Speed on a base below one just makes small things smaller, faster.

02 **Literacy as paperwork** Symptom: a policy on a shelf, a slide deck, a box ticked. Article 4 treated as a document instead of a behaviour.

THE CORRECTION

The test is behaviour, not a binder. Teach literacy on real work — what is safe to put in, how to brief, which calls a machine must not make alone — and keep the evidence.

03 **ROI theatre** Symptom: a launch event, a headline number, a victory lap — then the quiet rehire. The win is announced before it is real.

THE CORRECTION

Klarna declared victory, then rehired. Measure the boring lines a CFO can find, and hold satisfaction and quality flat. Cancel the win that only exists in the press release.

A stronger exponent on a base below one produces a worse result, faster.

THE LAW OF THE ROOF

Measure three things. Refuse to claim a fourth.

The protocol is judged on movement you can show, not on a number you wish were true. Three measures tell you whether the base is rising. A single universal ROI figure is not one of them.

01 / TIME

Hours reclaimed

The work about work you took back – 58% of the day is the ceiling. Count the hours freed, then where they went.

RECLAIMED · REDEPLOYED

02 / FLOW

Processes moved

How many tribal processes are now written, wired, and firing without a human chasing. Count the ones that became reflexes.

ENCODED · AUTOMATED

03 / PEOPLE

Rungs climbed

How many people moved from operator to orchestrator – directing the machine, owning the verdict, spending reclaimed hours on judgment.

OPERATOR → ORCHESTRATOR

NOT

Not a universal ROI figure. There is no honest single multiple that travels across companies – the return depends entirely on the base you applied the exponent to. Publish the rubric and your own measured lines; never claim a number that holds for everyone. The one figure we will stand behind is the base rate you are trying to beat: **95% returned zero.**

Print this page. One week a row. Tick it when it is true.

DAYS 1-30 · DIAGNOSE — FIND THE WEAKEST LAYER

W1 **Score the Stack.** Grade all five layers, true on a bad week. Mark the lowest red.

DAYS 1-7

W2 **Compute revenue per employee.** Write it beside the benchmarks (median ~\$130K; AI-native \$3-7M).

DAYS 8-14

W3 **Aim at one P&L line** with a date. **Cancel** every initiative that has none.

DAYS 15-21

W4 **Name the one 30-day move** for the red layer. Exit test: red layer, number and date all named.

DAYS 22-30

DAYS 31-60 · FIX THE BASE — RAISE IT ABOVE ONE

W5 **Clean one process.** Excavate input, action, output. Run the five-whys; delete the ghost steps.

DAYS 31-37

W6 **Free three data hostages.** Move each fact only one person holds into the system, not a head.

DAYS 38-44

W7 **Stand up one agent pilot** in policy: minimum tools, a human on anything consequential, a small eval.

DAYS 45-51

W8 **Wire one handoff** so the next step fires unchased. Exit test: one process firing, one source of truth.

DAYS 52-60

DAYS 61-90 · COMPOUND — TURN ON THE MULTIPLIER

W9 **Redeploy the reclaimed hours** to judgment and novelty — not to a smaller payroll.

DAYS 61-67

W10 **Hire one orchestrator,** not an operator: the person who designs and runs the system.

DAYS 68-74

W11 **Re-score the red layer.** Run the same audit. Record whether the number moved.

DAYS 75-81

W12 **Read the P&L line; set the next ninety days.** Name next quarter's red layer. Run it again.

DAYS 82-90

BLANKCOLLAR.UNIVERSITY — THE EXECUTION WING

Direct, don't do.

The base rate is settled. The method is the part that changes it. Run the ninety days in order — diagnose, fix the base, compound — and the same exponent that sank the 95% starts to pay. You do not need a bigger model. You need a base worth multiplying, and the discipline to build it first.

MOVE ONE

Run the equation.

Score your framework in about five minutes — a measured read that places you on the band and names the layer to fix first.

The Index · [the-index](#) · free

MOVE TWO

Book a training.

Run the ninety days with people who have done it — the programmes, formats and rubrics that turn this guide into a base above one.

Book a training · [/book](#)

Nobody is born into the five percent.
You choose your way in.

BLANKCOLLAR.UNIVERSITY
THEBLANKCOLLAR.COM

Every figure traces to a source. None is invented.

The numbers here are the report's numbers, kept straight. Where a figure is illustrative of a principle rather than a single measured study, it is marked as such.

-
- 01 **95% returned zero P&L impact.** MIT NANDA, 2025 — the measured profit-and-loss impact of corporate AI transformations at scale. The base rate the protocol is built to beat.
-
- 02 **78% adoption; 80%+ no earnings impact.** McKinsey, 2025 — AI used in at least one function by most companies, with the majority reporting no material impact on earnings.
-
- 03 **58% of the day on work about work.** Industry research cited in The Blank Collar — time lost to searching, chasing and confirming which number is real.
-
- 04 **~94% of routine HR tasks automated, total employment rose.** IBM workforce disclosure — the hours were redeployed, not eliminated.
-
- 05 **\$3-7M revenue per employee at AI-native firms; ~\$130K median software.** Cursor and comparison figures cited in The Blank Collar — leverage, not harder work.
-
- 06 **The three autopsies.** McDonald's × IBM (a Process failure), Air Canada (a governance / Human-Experience failure), Klarna (a Vision failure) — public cases read in The Blank Collar.
-
- 07 **The 57-step approval; the jagged frontier; the exponent.** Illustrative of the Process, Human-Experience and AI layers respectively, as set out in The Blank Collar (2nd edition).
-

This field guide is the operating companion to **The State of AI at Work — 2026**. It re-uses that report's findings and figures; it does not introduce new claims. A practical starting point, not a compliance opinion — confirm any regulatory obligations with qualified counsel.

90 / days

**Diagnose. Fix the base.
Compound. Then run it
again, every quarter.**

The 5% are not smarter or richer. They are the ones who fixed the base before they turned on the roof – and they ran a sequence you can start this quarter. Six moves, ninety days, one drawing.



Work is for bots. Life is for humans.

BLANKCOLLAR.UNIVERSITY
THE EXECUTION WING