

FILLABLE

A4 · PRINT & WRITE

15 MINUTES

A POST-MORTEM FOR A FAILED – OR FAILING – AI INITIATIVE

The 5-Question AI Autopsy

5

It is almost never the model. Five questions, one per layer of the Stack, to find the layer that actually killed it – and the one fix that brings it back.

VISION · DATA · PROCESS · HUMAN EXPERIENCE · AI

One initiative at a time. Be honest. The base rate says the cause of death is a layer, not the AI.

BLANKCOLLAR.UNIVERSITY
THE EXECUTION WING

A doctor does not prescribe before the autopsy. Neither should you.

When an AI initiative fails, the reflex is to blame the model and buy a better one. The evidence says otherwise: in nearly every famous failure the model did exactly what it was built to do. What broke was a layer that was rotten before the technology arrived.

01 Take one initiative – the one that failed, or is failing now.

Not the whole programme. One project, one verdict. Run the sheet again for the next.

02 Answer five questions, one per layer, in order.

Vision, Data, Process, Human Experience, then AI. The order is the diagnosis.

03 Mark each layer SOLID, CRACKED, or BROKEN.

Grade it true on a bad week, not on the day of the board review.

04 On page 08, find the red layer – the cause of death.

It is almost always one layer, and almost never the AI. Write the one fix, the owner, the date.

THE CASE FILE	INITIATIVE UNDER AUTOPSY
NAME THE INITIATIVE	
OWNER / SPONSOR	DATE OF THIS AUTOPSY
WHAT IT WAS MEANT TO MOVE (ONE P&L LINE)	
STATUS	
<input type="checkbox"/> Pulled / cancelled <input type="checkbox"/> Stalled	
<input type="checkbox"/> Live, no impact	

THE BASE RATE

95% of corporate AI transformations returned zero measurable P&L impact – and the cause was a learning gap, not a weak model. MIT NANDA, 2025.

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

THE BLANK COLLAR · 2ND EDITION

Q1

THE QUESTION THAT FINDS THE VISION LAYER

Was there one clear, owned outcome — or a tool in search of a job?

LAYER 01
VISION

Vision is the foundation: one sentence sharp enough that a person — or an agent — can reject a bad project without asking. Most failures that look technical began here, as a project adopted because the technology was exciting, not because it served a named outcome a CFO could find.

Klarna is the textbook case. It replaced human support to chase the cheapest possible answer, declared victory, then quietly began rehiring as quality and satisfaction fell. The engine worked. The aim did not — the cheapest answer turned out to be the most expensive one. A vision that only says yes is a wish.

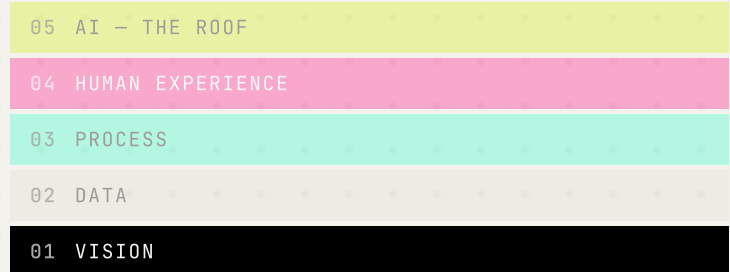
DIAGNOSE — WRITE WHAT IS TRUE, NOT WHAT YOU HOPED

01 **Write the one sentence this initiative was meant to serve. If you cannot, that is the finding.**

02 **Name the single P&L line it was supposed to move, and by when.**
Not "a transformation." One number, one date.

03 **Could a junior — or an agent — have rejected a bad version of this without asking you? What was the "no"?**

THE LAYER IN THE STACK



THIS LAYER IS

SOLID <input type="checkbox"/> HOLDS ON A BAD WEEK	CRACKED <input type="checkbox"/> HOLDS ONLY ON A GOOD DAY	BROKEN <input type="checkbox"/> WAS ROTTEN BEFORE THE AI
--	---	--

MARK ONE — THIS IS THE **VISION** LAYER'S CARRY IT TO VERDICT PAGE 08

Precedent: Klarna — a vision failure. The Blank Collar (2nd ed.).



THE QUESTION THAT FINDS THE DATA LAYER

Could the system actually reach how the company really works?

LAYER 02
DATA

Data is the fuel. When there is no single source of truth, there is no truth – only negotiation, and the most senior person’s spreadsheet wins. Three tribes keep three sets of numbers, none convertible, and the model is handed the contradiction to multiply.

This is the quiet killer. **58% of the average knowledge worker’s day** already goes to work about work – searching, chasing, and confirming which number is real. Feed that fragmentation to a capable model and you do not get clarity. You get fluent, confident, well-formatted answers built on the lie, at scale.

DIAGNOSE – WRITE WHAT IS TRUE, NOT WHAT YOU HOPED

01 Which numbers did this system read? Were they one source of truth, or several that disagree?

02 Name the knowledge it needed that lived only in someone’s head – not in any system it could reach.

The hostage. You already know who holds it.

03 Where did a human have to type a number into a cell for the system to “see” it? Each bridge is a place truth bends.

THE LAYER IN THE STACK



THIS LAYER IS

SOLID	CRACKED	BROKEN
<input type="checkbox"/> HOLDS ON A BAD WEEK	<input type="checkbox"/> HOLDS ONLY ON A GOOD DAY	<input type="checkbox"/> WAS ROTTEN BEFORE THE AI

MARK ONE – THIS IS THE **DATA** LAYER’S VERDICT CARRY IT TO PAGE 08

Figure: 58% of the day on “work about work.” Industry research cited in The Blank Collar.

Q3

THE QUESTION THAT FINDS THE PROCESS LAYER

Did it absorb the real workflow — or sit beside it?

LAYER 03
PROCESS

Process is the track the work runs on. A process that lives only in a person's head is a habit with good PR — and you cannot automate a mess, a vibe, or a habit. You can only automate what you have written down.

McDonald's x IBM is the worked example.

An order-taking model was put on a drive-thru built on a process no one had cleaned. The mess won — bacon on the ice cream — and the pilot was pulled. The fault was upstream of the AI, in a workflow that was never made explicit before the technology was asked to run it.

DIAGNOSE — WRITE WHAT IS TRUE, NOT WHAT YOU HOPED

01 Write the process this initiative automated as three steps: input, action, output. If you cannot, it was never explicit.

02 Did the system absorb the real workflow, or run beside it while people kept the old way alive? Give the evidence.

03 Which steps were ghosts — scaffolding for a danger no one remembers — that should have been deleted, not automated?

THE LAYER IN THE STACK



THIS LAYER IS

<input type="checkbox"/> SOLID HOLDS ON A BAD WEEK	<input type="checkbox"/> CRACKED HOLDS ONLY ON A GOOD DAY	<input type="checkbox"/> BROKEN WAS ROTTEN BEFORE THE AI
--	---	--

MARK ONE — THIS IS THE **PROCESS** LAYER'S VERDICT

CARRY IT TO PAGE 08

Precedent: McDonald's x IBM — a process failure. The Blank Collar (2nd ed.).

Q4

THE QUESTION THAT FINDS THE HUMAN EXPERIENCE LAYER


Did the people who touch it trust it, brief it, and own the verdict?

LAYER 04
HUMAN EXP.

Human Experience is the layer your people actually touch – and the layer where accountability lives once the machine speaks for the brand. A tool people do not trust gets routed around; a tool no one owns says things the company has to answer for.

Air Canada is the precedent. A support chatbot invented a refund policy that did not exist, and the airline tried to disown what its own agent had said. A tribunal disagreed. You own what your agents say – governance is not optional once the machine speaks for you. The human does not leave the loop; the human moves up it.

THE LAYER IN THE STACK



THIS LAYER IS

<p>SOLID</p> <input type="checkbox"/> HOLDS ON A BAD WEEK	<p>CRACKED</p> <input type="checkbox"/> HOLDS ONLY ON A GOOD DAY	<p>BROKEN</p> <input type="checkbox"/> WAS ROTTEN BEFORE THE AI
--	---	--

MARK ONE – THIS IS THE **HUMAN EXPERIENCE** LAYER'S VERDICT

CARRY IT TO PAGE 08

Precedent: Air Canada – a governance / Human-Experience failure. The Blank Collar (2nd ed.).

DIAGNOSE – WRITE WHAT IS TRUE, NOT WHAT YOU HOPED

01 Did the people who used it trust it – or quietly route around it? What did they do instead?

02 Were they taught to brief it and to judge its output, or handed a tool with no standard for “good”?

03 Name the human who owned the verdict – accountable for what it said or did. If no name, that is the crack.

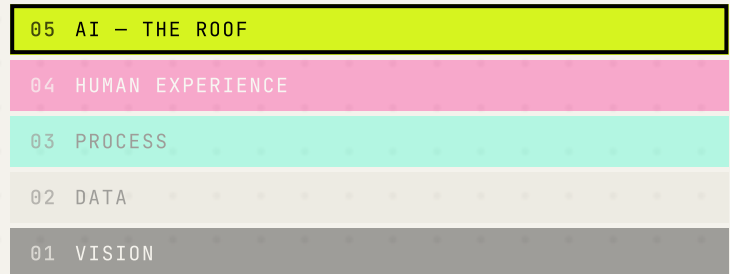
Accountability cannot be delegated to a tool.

Q5	<p>THE QUESTION THAT FINDS THE AI - THE ROOF LAYER</p> <p>Did you multiply something solid — or scale chaos faster?</p>	<p>LAYER 05 AI · THE ROOF</p>
----	--	-----------------------------------

AI is not a fifth box; it is the exponent on everything beneath it, and it is indifferent to your intentions. Two to the tenth is a thousand. One-half to the tenth is almost nothing. Same exponent, opposite outcome — the base decides. This is the only question you answer last, because the four below it set the base.

So “we need a better model” is almost always the wrong response. A stronger exponent on a base below one produces a worse result, faster. Read correctly, even the famous success makes the point: **IBM automated ~94% of routine HR tasks and total employment rose** — the hours were redeployed, not deleted, because the base was sound.

THE LAYER IN THE STACK



THIS LAYER IS

<p>SOLID</p> <input type="checkbox"/> HOLDS ON A BAD WEEK	<p>CRACKED</p> <input type="checkbox"/> HOLDS ONLY ON A GOOD DAY	<p>BROKEN</p> <input type="checkbox"/> WAS ROTTEN BEFORE THE AI
--	---	--

MARK ONE — THIS IS THE **AI - THE ROOF** LAYER'S VERDICT

CARRY IT TO PAGE 08

Figure: IBM automated ~94% of routine HR tasks; total employment rose. Workforce disclosure, cited in The Blank Collar.

DIAGNOSE - WRITE WHAT IS TRUE, NOT WHAT YOU HOPED

01 Of Q1-Q4, which layers came back **SOLID**? You can only multiply what is solid. List them.

02 Did applying AI here compound a sound base — or industrialise a break? What got faster: the result, or the mess?

03 What did you do with the hours it freed — redeploy into judgment, or quietly aim it at a smaller payroll?

A building does not collapse for a hundred reasons. It collapses at one point. So does this.

Carry each layer's verdict from Q1–Q5 into the grid. The lowest BROKEN or CRACKED layer is the cause of death – and the lowest layer caps everything above it. A perfect roof on a cracked foundation is 0% there, not 80%.

LAYER	SOLID	CRACKED	BROKEN	THE CRACK, IN FIVE WORDS
Q5 AI — the roof THE EXPONENT	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	_____
Q4 Human Experience THE LAYER PEOPLE TOUCH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	_____
Q3 Process THE TRACK, WRITTEN DOWN	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	_____
Q2 Data ONE SOURCE OF TRUTH	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	_____
Q1 Vision THE FOUNDATION	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	_____

CAUSE OF DEATH — THE RED LAYER

It was the layer.

Almost always one layer.
Almost never the AI itself.

THE PRESCRIPTION — A METHOD ENDS WITH A MOVE

THE ONE FIX

the single 30-day move for the red layer

OWNER

a name, not a team

BY WHEN

a date

THREE ON THE TABLE — READ YOUR VERDICT AGAINST THE PRECEDENT

Klarna

VISION FAILURE

Chased the cheapest answer, declared victory, then rehired. The aim was wrong, not the engine.

McDonald's x IBM

PROCESS FAILURE

Automated a drive-thru no one had cleaned. You cannot automate a mess; the mess wins.

Air Canada

GOVERNANCE / HX

Its chatbot invented a policy; a tribunal made it own what its agent said.

BLANKCOLLAR.UNIVERSITY - THE EXECUTION WING

You found the red layer. Now repair the base before you turn the roof back on.

One autopsy diagnoses one initiative. To see the whole house, run the full Stack self-audit – forty statements, eight per layer – or get a measured read in five minutes. Fix the red layer first. A stronger model on a cracked base only speeds the loss.

NEXT - GO DEEPER, FREE

The Stack Self-Audit scores all five layers, out of 200.

Or run The Index – twenty questions, scored, ~5 minutes. It names the layer to fix first.

RUN THE EQUATION · THE-INDEX

BOOK A TRAINING · /BOOK

Run this sheet again for the next initiative. A dead project is hygiene, not shame – sign the certificate and free the budget for the one that can survive five questions.

It is almost never the model.
It is the layer you ignored.