

The Remediation That Cost More

Arc 1 · Post-mortem / write-down

In the room: Whoever has to explain it — ops partner, IC, LP advisory

The question: What did the fix cost, and why didn't we stop it?

THE POSITION

SIGHTING

The most expensive automation failure isn't the one that doesn't work. It's the remediation — aimed at the part that failed, landing on the part that was working.

THE CONSENSUS READ

A failed build. Write off the capex and move on.

THE MECHANISM

The write-off is the small number. A fix operators stop trusting doesn't get escalated — they quietly reinsert themselves and check everything, erasing the gain on the ninety percent that never had a problem. That regression is cycle time drifting back to baseline with no line item attached. And the person who must kill it approved it, which is measurably the worst possible evaluator.

THE EXPOSURE

The reported number understates the damage by roughly 3.5 to 4.4 times.

THE TEST

Find a remediation approved in the last two years and pull the process metric before the original build, after it, and after the fix. If the third number is worse than the second, you have this.

THE BRIDGE — CLAIMED TO REALISED

The operating case is the same composite as Episode 01. The Dietvorst and Staw findings are real, published and cited with their samples; the two are kept separate throughout.

LINE	MOVES	RUNNING
Claimed — A \$180,000 remediation to recover \$260,000 of leaked automation savings, presented to the board as a nine-month payback		+\$80,000 net claimed
1. The build cost, which is the number everyone tracks	-\$180,000	-\$180,000
Six weeks of engineering, roughly \$180,000 all-in against the composite.		
2. The regression on the part that was working	-\$310,000 to -\$420,000	-\$490,000 to -\$600,000
The cost nobody books.		
3. Why it ran for eleven months after it was known to have failed	-\$140,000 to -\$190,000	-\$630,000 to -\$790,000
The escalation cost, and this line belongs to the sponsor rather than the portfolio company.		
4. Where I was wrong, twice — and the second one is the useful one [where I was wrong]	-\$0	-\$630,000 to -\$790,000
Both on the record because they compound. Cost: Staw's finding is that the binding variable is who reads the criterion, not how well the criterion is written.		
5. Total against \$260,000 recovered [residual]	+\$0	-\$630,000 to -\$790,000
Presented as a \$180,000 spend to recover \$260,000.		
Realised		-\$630,000 to -\$790,000 against a claimed +\$80,000 net

THE DOWNSIDE

Multiple: 9x, illustrative

\$630,000–\$790,000 of destroyed run-rate performance at an illustrative 9x, against a project whose written-off capex was \$180,000. Substitute your own comps.

\$5.7M to \$7.1M of enterprise value — against a reported write-off of \$180,000

When it surfaces. The capex write-off surfaces immediately and is discussed. The regression surfaces as a slow drift in an operating metric with no owner and no project attached — typically noticed at the next VCP review or, worse, by a buyer comparing the operating metrics in the CIM against the ones two years earlier.

The Remediation That Cost More

The question: What did the fix cost, and why didn't we stop it?
Paste into a request list or a management agenda. Each question resolves to an artifact, not to a characterisation.

THE DILIGENCE PACK

1. For each automated process with a remediation, provide the primary operating metric at three points: before the original build, after it, and after the remediation.

Artifact: Transaction-level data at all three points.
The third number against the second is the whole diagnosis.

2. Measure the manual check rate on the clean majority of transactions — the ones that never had a defect — before and after the remediation.

Artifact: Sampled observation, not self-report.
Nobody reports that they've started checking everything.

3. Identify who approved the original build, who approved the remediation, and who presented each review.

Artifact: The approval and board records.
If those are the same person, the review has a known and measured bias and should be read accordingly.

4. For any remediation currently running past its stated payback, list every additional analysis, pilot or data-gathering exercise commissioned since the payback date passed.

Artifact: The list.
Per Staw, post-threshold information-gathering is the escalation response. Treat it as a diagnostic, not as diligence.

5. State, in advance, the metric and threshold at which the next remediation stops, and name the person who will present that reading — who may not be the approver.

Artifact: A written line in the approval memo, before funding.
The one control that matters, and it is free.

DISQUALIFIER

If the answer to question 3 is that the approver also presented the review, you do not need the rest of the pack to know how to read the review. Before approving any fix, establish what is currently working and price the cost of breaking it — remediations are scoped against the failure and evaluated against the failure, while the exposure sits on the volume that never had a problem.

SOURCES — EVERY MEASUREMENT WITH ITS SAMPLE

CLAIM	SOURCE	SAMPLE / CLASS
Model outperformed humans in every condition; humans produced 15–29% more error on one task and 90–97% more on another. Participants who watched the model err were significantly less likely to bet on it, including the 83% who watched it beat the human. Seeing a human err did not reduce willingness to use the human.	Dietvorst, Simmons & Massey, Journal of Experimental Psychology: General 144(1), 2015	Five studies, incentivised choice with money at stake (measured)
~75% of subjects responsible for a prior failure sought retrospective justifying information, against ~25% of those not responsible	Staw, escalating commitment literature (1976 onward)	Experimental subjects; responsibility manipulated (measured)
\$180,000 build at ~61% classifier precision; regression across the clean 90% of volume; eleven months of duplicated effort	Composite, built to the shape of a real quote-to-cash automation (same case as Episode 01)	n/a — composite (composite)
9x entry multiple	Illustrative mid-market comp	n/a — illustrative, arithmetic exposed (illustrative)

One standing caveat. Every number in this show is somebody else's measurement, and I'll tell you whose, with the sample. None of it is diligence on your deal. Do that yourself.