RJL SOLUTIONS

EXECUTIVE PERSPECTIVE / CHANGE MANAGEMENT

Adoption by Design

The change management that decides whether enterprise AI is used.

OWNERS · HUBS · EVERYDAY WORK

Richard White
RJL Solutions · rjlsol.com

OCTOBER 2026

01 / THE PATTERN

A system that works can still go unused.

Most enterprise AI programs plan carefully for the model, the data and the controls. Far fewer plan for the people whose work will change. Skip that, and the technology can pass every test and still fail to change the business.

This paper is for leaders who already have a working capability, or soon will. It covers the change management that sits beside the technical foundation in the other RJL papers: who owns the change, how teams organize around it and how a hackathon can start something that lasts.

What the evidence says, and how far to trust it

SOURCE	REPORTED FINDING	HOW TO READ IT
RAND, 2024 [1]	Leading causes of failure included leaders misunderstanding the problem to solve, and teams chasing new technology over user needs	Interviews with 65 experienced practitioners, not a measured failure rate
BCG, 2026 [2]	About 10 percent of AI value comes from algorithms, 20 from technology and 70 from rethinking people and how they work	A heuristic from consulting case work, not a measured law
MIT NANDA, 2025 [3]	About 95 percent of organizations reported no measurable return on generative AI, with workflow integration and learning cited	Preliminary findings from 52 interviews, 153 leader surveys and a review of 300+ public initiatives; the authors call them directional
Writer survey, 2025 [4]	31 percent of employees said they undermine their company’s AI strategy; 71 percent of executives said AI was built in silos	Vendor-sponsored survey of 1,600 U.S. knowledge workers

Don’t quote any of these figures as the failure rate of your program. Taken together, though, they point one way. The hard part is rarely the model.

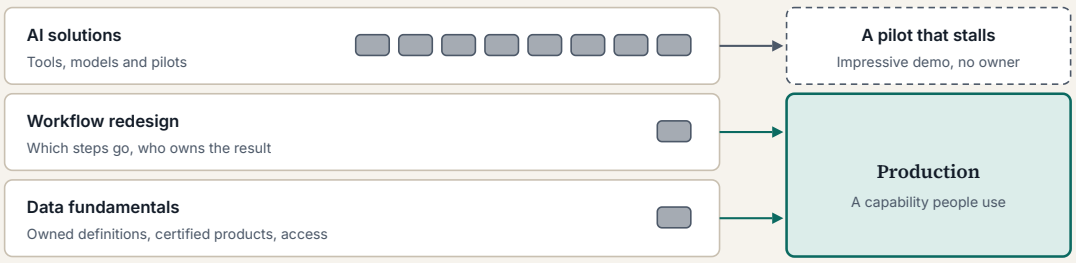


Exhibit: Everyone queues for AI solutions; the open lanes lead to production.

THE TEST Ask who will do their job differently on the first Monday after release, and whether they have agreed to.

02 / OWNERSHIP

Name the person who changes the work, not only the person who funds it.

A sponsor approves the budget. An owner runs the workflow and answers for its result. A change leader helps people move. One person can hold two of these roles. A program with only a sponsor, though, has nobody accountable for use.

Begin with the reason, tied to a measure

Sol Rashidi urges leaders to ask why before choosing any technology, and to tie conviction to business outcomes instead of enthusiasm for AI [5]. A useful answer names something the business already measures: cycle time, rework, error rate or cost per case. Enthusiasm isn't an outcome. It won't move an organization that's comfortable as it is.

Align the top team before the work starts

Executives who disagree about purpose hand the disagreement down as ambiguity. Before a proof starts, get the sponsor, the owner, the technology and risk leaders and the leader of the affected function to agree one page: the workflow, the outcome, what changes for whom and what will stop.

ROLE	ACCOUNTABLE FOR	EVIDENCE IT IS REAL
Sponsor	Funding, priority and clearing blockers	Attends the decision review, and has stopped something before
Business owner	The workflow, the outcome measure and acceptance	Can change the process without asking another team
Change leader	Communication, training and the feedback loop	Has protected time, and users know their name
Affected function leader	Roles, incentives and team capacity	Has agreed what the team will do differently

CONVICTION NEEDS A NUMBER

Moving a comfortable organization takes sponsorship tied to a measured outcome. A sponsor who can't say which number will move hasn't yet decided why.

03 / PATTERN ONE

Consolidating teams does not integrate their work.

Here's a familiar pattern in large institutions. Separate analytics groups, each serving different decision-makers, get brought together under one mandate for sound reasons: shared data, shared methods, a single view across several functions. The organization chart changes. The work doesn't.

This is a composite drawn from my own experience in regulated institutions, with details changed. Each group kept its own data, its own methods and its own relationships. Sharing meant giving up control of work people had built and been recognized for. Nobody had a reason to do it. The combined unit drifted further from the business it was meant to serve. People stayed where they were comfortable, and many didn't want the arrangement to work. Leaders talked about integration while the daily work stayed separate.

Why it happens

WHAT LEADERSHIP ASSUMED	WHAT PEOPLE EXPERIENCED
Shared data would pull teams together	Sharing meant losing control of data they had built and been recognized for
A common mandate would create common work	Each group kept its own deliverables and its own customers
Distance from the business was an efficiency	Fewer conversations with users meant less reason to change
People would value the larger mission	The comfortable position was the current one

What to do instead

Consolidate around shared work. Shared reporting lines won't do it. Choose two or three products several groups need, name one owner for each and make their use measurable. Put a business co-lead beside the technical lead. Rotate people between the groups and the business. And recognize the shared product, not only each group's own output.

WARNING SIGN — If anyone describes integration only in terms of reporting lines and tools, the work hasn't been integrated.

04 / PATTERN TWO

Experts will judge the system against their own work.

The second pattern shows up when AI reaches skilled professionals. Their judgment is the benchmark. Their role is what's at risk. And they're often the people best placed to find what the system gets wrong.

In this composite, a team built an AI assistant for experienced users and pitched it as a way to save them time. The users tested it against what they knew they could do themselves. They found the edge cases where it was weaker, and concluded they could do the work better. Some of that was accurate. Some of it was also a defense of roles they didn't want diminished. The project treated them as recipients of a tool. It never made them its designers, and it never said clearly what their role would be afterward.

Treat the objection as information

Two things can be true at once. The expert may be right about a real flaw, and may also be protecting a position. Test the first. Address the second openly.

WHAT THE EXPERT SAYS	WHAT IT MAY MEAN	WHAT TO DO
I can do this better myself	A real edge case, or a role under threat	Run their cases, publish the results, and define where their judgment is the control
It is not accurate enough	A fair standard that was never agreed	Agree on acceptance criteria with them before the next release
It will take my job	A fear no one has answered	State the future role plainly and early, and promise only what you can keep
Nobody asked us	Exclusion from the design	Make them co-designers and evaluators

What research on expert judgment suggests

Experiments by Dietvorst, Simmons and Massey found that people lose confidence in an algorithm faster than in a human after seeing the same mistake [6]. A follow-up found people far more willing to use an imperfect algorithm when they could adjust its output, even slightly [7]. The design lesson: give experts a bounded way to correct the system, record each correction and use the record to improve it. Reviews of clinical alerts show users override automated prompts that are irrelevant or too frequent [8].

DESIGN PRINCIPLE — Give the experts authority over the evaluation. People who wrote the test can't easily dismiss the result.

05 / PUBLIC CASES

What public evidence shows about people and programs.

Public evidence is thinner than anecdote suggests, and it's mixed. I've included these cases because their causes are documented, and I read each one with its limits.

CASE	WHAT IS DOCUMENTED	WHAT IT SUGGESTS	LIMITS
NHS National Programme for IT, 2000s [9]	MPs reviewing the program reported that clinicians were poorly consulted in the early design and were skeptical, alongside supplier delays	A program designed for users without them loses their confidence	A review of a very large program with many causes, not a study of adoption alone
MD Anderson and IBM Watson, to 2016 [10]	A University of Texas audit found about $62 million spent largely outside normal procurement. The system never went live and was not used for patient treatment apart from testing	Ambitious AI without ownership and normal governance can spend heavily and still not reach users	The audit examined procurement and governance, not the science or clinician views
Hackathon projects, 2020 [11]	Of 592 projects, about a third showed any activity afterward, and about 3.5 percent five months later	Energy at an event does not become a program without owners and follow-on	Public events and code repositories, not corporate programs
Enterprise AI surveys, 2025 [3] [4]	Organizations reporting no measurable return, shadow use of personal tools, and employees who say they undermine AI plans	Adoption is a management problem as well as a technical one	Vendor-linked or non-random samples; read as direction, not rates

The two composites in this paper, both from my own experience, point the same way. In each, the technology was acceptable. What was missing was a shared owner, a reason for the affected people to change and a way for them to shape the result.

READING THE EVIDENCE Use these cases to ask better questions of your own program. Don't use them to predict its outcome, and never cite the percentages as failure rates.

06 / RESISTANCE

Map the people before you map the plan.

Resistance is rarely irrational. People protect their skills, status, time and judgment. It also rarely announces itself. You see it as low repeat use, workarounds and personal tools used outside the approved route.

Engagement failures are well documented in large programs. MPs reviewing the NHS National Programme for IT heard that clinicians had been poorly consulted in the early design and were skeptical of it, among other problems [9]. Rashidi describes ten AI archetypes, including the Naysayer, the Evangelist and the Doer [5]. The table below is our own shorthand. It isn't there to label colleagues. It's there so you expect different motives and prepare for each.

STAKEHOLDER PATTERN	TYPICAL MOTIVE	WHAT HELPS
Skeptic of the outcome	Has seen earlier tools fail	Show evidence on their own cases
Enthusiast	Wants visible wins, may overpromise	Give a bounded role and hold them to the measure
Quiet expert	Fears loss of status or autonomy	Make them the evaluator
Gatekeeper	Controls data, systems or approval	Offer shared ownership and shared credit
Overloaded operator	Has no capacity to change	Remove work before adding work

In the Writer survey, 31 percent of employees said they undermine their company's AI strategy, citing fear of AI and concern about tool quality [4]. Both have answers. Fear needs a plain account of the future role. Quality concerns need a test the skeptic helped write.

PRACTICE Hold a short listening round with the affected team before design. Record each concern in a log with an owner and a date for the response.

07 / THE WORK

Change the process, the roles and the incentives, not only the tool.

Add AI to an unchanged process and you mostly add a step. Both the MIT and BCG reports point to workflow and people, not model quality, as where value gets lost [2][3].

Redesign four things together

ELEMENT	QUESTION TO SETTLE BEFORE RELEASE
Process	Which steps change, stop or move, and where does human review sit
Roles	What the person does now: judgment, exceptions and oversight, and who answers for the result
Incentives	Are people measured and rewarded for what the new process needs
Capacity	What work is removed so people have time to learn and to supervise

Decide where AI should not act

Not every task should go to AI to the same degree. Two questions decide it: what would an error cost, and how much human judgment does the task need? Each quadrant carries a control, agreed with the owner before release.

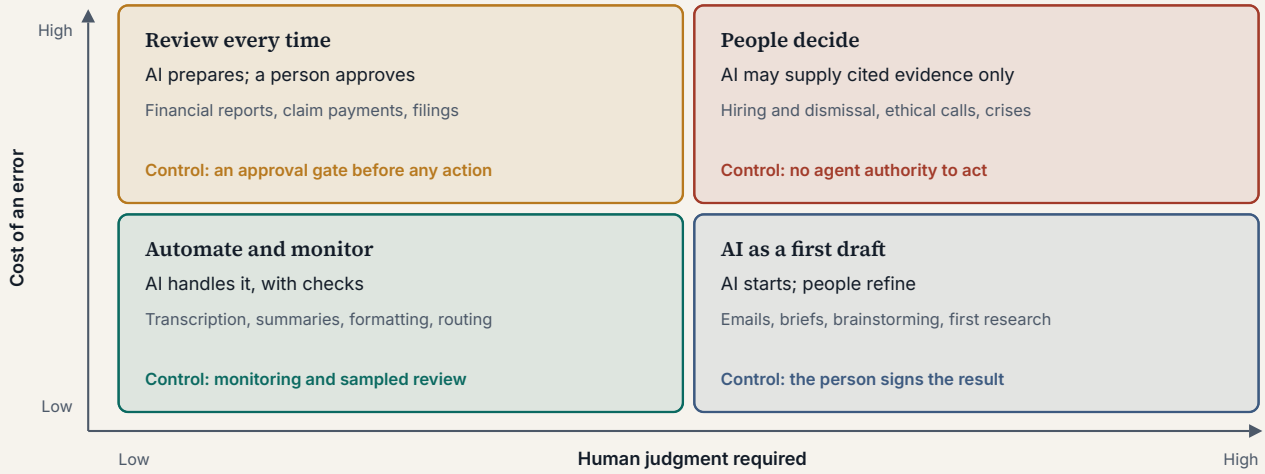


Exhibit: The cost of an error and the judgment required decide how much AI may do.

Design oversight, and the reviewer's job, deliberately.

Rashidi argues that oversight should be designed into the process from the start, with accountability named, instead of added after something goes wrong [5]. Oversight has a cost. In a 2026 interview she said she had switched off two of her four AI agents, because supervising them took more time than the work they returned [12].

That's why the RJL diagnostic asks for cost per completed task, including review and rework, and not just the cost of a model call. A capability that needs more supervision than it saves will be abandoned, whether or not anyone announces it.

Make review a designed job

Human review protects the work only when someone designs it as a job. Size the review capacity the volume will need before launch, not after the queue forms. Give each approver the information, the time and the authority to reject or change the result. Watch for automation bias, the tendency to accept a fluent recommendation without checking it; the EU AI Act names it among the risks human oversight of high-risk systems must address [15]. Track override rates and review times. An approver who never disagrees may not be reviewing.

DECISION If supervising the system costs more than the work it saves, change the design or stop. Agree that rule with the owner before launch.

08 / HUBS

Organize around hubs where business and technology share one backlog.

A hub is a standing group built around one business domain, such as claims, underwriting, finance or risk. It brings business and technology people together around one backlog, one set of data products and one outcome measure. It's a way of working across existing departments. It isn't a new department.

HUB ELEMENT	PRACTICE
Leadership	A business lead and a technology lead, jointly accountable for the outcome
Membership	Business experts, a data product owner, an engineer and a risk partner, some part time
Backlog	One ranked list of workflows, owned by the business lead
Shared assets	Certified data products and approved definitions that the hub maintains
Connection	A small central platform group provides tools, standards and controls to every hub

Why hubs counter the silo pattern

Members share work on real problems instead of sharing reports. The business sees progress and has a say in priorities. The central group sets standards and runs the platform without owning every use case. The hub is also where the people who depend on certified data products, as the Data Product Certification Standard describes them, use and maintain them.

Start with one or two

Begin with the domains that have the clearest outcome and an owner who wants the work. Learn, then replicate. A hub that drifts into another consolidated team repeats pattern one. Test it with two questions. Do the business members spend time with users? Does the backlog change decisions?

TEST A hub is working when the business lead can name the last three things it changed in daily work.

09 / HACKATHONS

Use hackathons to start hubs, not to produce demos.

A hackathon that brings business and technology people together for two or three days can close distance, surface real problems and build working relationships fast. On its own, it rarely produces anything that lasts.

Nolte, Chounta and Herbsleb studied 592 hackathon projects. About a third showed any activity after the event, and about 3.5 percent were still active five months later [11]. Their data came from public events and code repositories, not corporate programs, so treat it as a caution, not a forecast. Projects that continued tended to have diverse skills in the team, a good match between skills and project, and an intention to reach real users.

Design for the day after

STAGE	PRACTICE
Before	Choose problems with a named business owner and a baseline measure. Pre-approve governed sandbox data. Confirm the follow-on funding and decision date.
During	Mixed teams of business experts, engineers, a data owner and a risk partner, working a real workflow. Users test on day two.
At the end	Judge on three tests: does it solve a measured problem, would users adopt it, and could it pass certification and release gates.
After	Selected ideas receive two to four weeks of funded follow-on, a named owner, and a decision to continue in Prove, hold or stop.

A hackathon prototype isn’t production. It still has to pass data certification, access control and release evidence like any other release. Reward adoption and outcomes. Don’t reward demonstrations.

RULE Every hackathon project ends with a decision and an owner, or it ends.

Train for the new work, and measure whether it is used.

Training on a tool isn't training for a job. People need to know when to trust the system, when to challenge it and when to escalate. Managers need to know how to use it in decisions.

BCG reports that the companies it calls future-built plan to upskill more than half their employees, against about a fifth of laggards, and protect time for learning [2]. Article 4 of the EU AI Act, as amended, asks providers and deployers to take measures that support AI literacy among their staff [13]. It's a duty to act, not to guarantee a level. It's also a reasonable floor for any regulated enterprise.

Measure adoption as an outcome

MEASURE	WHAT IT TELLS YOU
Eligible users and repeat use	Whether the workflow is part of the work
Override and escalation rates	Where users distrust the system or it is weak
Outcome against baseline	Whether the business result moved
Review and rework time	The true operating cost
Workarounds and personal tools	Where the approved route is failing
A short user pulse	An early warning, not proof

Gartner expects more than 40 percent of agentic AI projects to be canceled by the end of 2027, citing cost, unclear value and weak risk controls [14]. A baseline and an accountable owner let you show value, or stop early.

BASELINE FIRST Record the baseline before launch. Without it you can't show success, and you can't stop a failing capability early on evidence.

Add the people work to each phase.

Change management isn't a separate stream running beside delivery. It belongs in each phase of Frame, Prove, Build and Run, with its own evidence for the decision at the end.

PHASE	PEOPLE WORK	EVIDENCE FOR THE DECISION
Frame	Align the sponsor, owner and affected leaders. Map stakeholders. Define what changes for whom. Choose the workflow with its owner.	A signed one-page alignment, a named owner and a named change leader
Prove	Put real users in the slice as evaluators. Run a mixed hackathon to select ideas. Test the acceptance criteria with the experts.	Users' evaluation of their own cases, and their stated intent to use it
Build	Redesign process, roles and incentives. Train. Communicate. Stand up the hub.	Training complete, roles agreed, oversight capacity funded
Run & Transfer	Measure adoption and outcome. Keep the feedback log. Name a successor owner. Refresh training.	Adoption and outcome reviewed with the owner, and a successor named

Where to begin this month

Take one workflow that's already live, or close to it. Ask the owner the eight questions on the next page. If any answer is No or Unknown, hold a listening round with the affected team and agree the one-page alignment before the next release. Then use the online diagnostic for the technical questions, and bring both to a working session.

WHAT A DECISION NEEDS	Evidence that the people whose work will change have agreed to it, and have what they need to do it. A launch date and a trained model don't count.

12 / WORKSHEET

Adoption readiness: eight questions to answer with the owner and the affected team.

Answer from what's been done. The plan doesn't count. For each question, mark Yes, Partial, No or Unknown, and note the evidence an outsider could inspect. Never mark Yes for something announced but not tested with the people affected.

	A YES MEANS	EVIDENCE TO SHOW
A1	A named business owner has agreed to change the workflow and accept the result	Owner's decision rights and outcome measure
A2	The leader of each affected function has agreed what their team will do differently	The one-page alignment, signed
A3	The people affected were consulted, and their future role has been described to them	Listening-round log and the role statement
A4	Experts helped write the acceptance criteria and have tested their own cases	Criteria and their results
A5	Process, roles and incentives have been redesigned, not only the tool	Before and after process, and revised measures
A6	Training and oversight capacity are funded and scheduled	Plan, budget and named supervisors
A7	A hub or joint backlog connects the business and technology	Backlog, members and lead names
A8	A pre-launch baseline exists, and adoption is measured against it	Baseline, adoption and outcome measures

Count the answers. Don't turn the count into a score. A single critical gap, such as no agreement from the affected leader, can outweigh several strong answers. The most useful result is usually a question on which the owner, the technical lead and the affected team disagree.

THEN Turn the three gaps that matter most into actions with an owner and a date. Revisit them at the next release review.

Sources and how they were used.

Prepared by Richard White, 9 October 2026. The two patterns in this paper are composites from the author's experience in large institutions, with details changed. The sources below are public. Survey and interview figures are reported as the sources state them, with their limits.

[1] RAND Corporation, Ryseff, J., De Bruhl, B. F. and Newberry, S. J., *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed*, RR-A2680-1, August 2024. Interviews with 65 experienced data scientists and engineers.

[2] Boston Consulting Group, Bedard, J. and Beauchene, V., *AI Transformation Is a Workforce Transformation*, 4 February 2026. The 10-20-70 principle is drawn from BCG case work.

[3] Challapally, A., Pease, C., Raskar, R. and Chari, P., *The GenAI Divide: State of AI in Business 2025*, MIT NANDA, July 2025. Preliminary findings from a review of over 300 public initiatives, 52 organization interviews and 153 senior-leader surveys.

[4] Writer and Workplace Intelligence, "68% of C-suite say AI adoption has caused division at their company," press release, 18 March 2025. Survey of 1,600 U.S. knowledge workers (800 executives, 800 employees), fielded November and December 2024. Vendor-sponsored.

[5] Rashidi, S., *Your AI Survival Guide*, Wiley, 2024, drawing on chapters 1, 3, 4, 5 and 6. Her ideas are paraphrased here, not quoted.

[6] Dietvorst, B. J., Simmons, J. P. and Massey, C., "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General* 144(1), 2015, 114–126.

[7] Dietvorst, B. J., Simmons, J. P. and Massey, C., "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science* 64(3), 2018, 1155–1170.

[8] van der Sijs, H., Aarts, J., Vulto, A. and Berg, M., "Overriding of Drug Safety Alerts in Computerized Physician Order Entry," *Journal of the American Medical Informatics Association* 13(2), 2006, 138–147.

[9] Cross, M., "Benefits of £12bn IT programme in NHS are 'unclear,' MPs say," *BMJ* 2007;334:815, 19 April 2007.

[10] University of Texas System Audit Office, special review of the MD Anderson Cancer Center Oncology Expert Advisor project, 2016, as reported by The Register, 20 February 2017.

[11] Nolte, A., Chounta, I.-A. and Herbsleb, J. D., "What Happens to All These Hackathon Projects? Identifying Factors to Promote Hackathon Project Continuation," *Proceedings of the ACM on Human-Computer Interaction* 4 (CSCW2), Article 145, 2020. doi:10.1145/3415216. 592 projects from public events.

[12] Business Insider interview with Sol Rashidi on shutting down AI agents, June 2026, as reported by Let's Data Science, 1 July 2026.

[13] Regulation (EU) 2024/1689 (the AI Act), Article 4, as amended by Regulation (EU) 2026/1744.

[14] Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," press release, 25 June 2025.

[15] Regulation (EU) 2024/1689 (the AI Act), Article 14, Human oversight, including the awareness of automation bias in Article 14(4)(b).

Proposed operating guidance. Enterprise policy and the named decision holders determine approval. Not legal advice.

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, and Run & Transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

PERSPECTIVE **Why Enterprise AI Stalls After the Demo**
The executive case for governed production, in ten pages.

WORKSHEET **The 10-Question AI Production-Readiness Diagnostic**
Four pages to complete with your business owner, technical lead and control owner.

GUIDE **The 90-Day First Governed Capability**
A delivery plan from Frame to an evidence-based investment decision.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

This paper describes no client engagement and no client result. The two patterns are composites. The worksheet is a discussion aid, not a certification. Survey figures appear as the sources state them; they aren't failure rates for any one program.