

You do not need AGI. You need agreement.

The hard part of enterprise AI is not getting a model into production. It is getting an organization to agree on what it is for, who owns the consequences, and how people work when the answer is contested.

WE DON'T NEED MORE INTELLIGENCE

We need organizations that can make a clear decision together, act on it, and keep that decision intact when the model, market, regulator, or leadership changes. **That is a leadership problem before it is a technology problem.**

Most companies are treating AI adoption as a tooling exercise while employees experience it as a mandate, a threat, or one more change happening to them. Often the resistance is not to the technology. It is to

absorbing ambiguity, risk, and broken process without understanding the intent or having any meaningful part in the decision. **That is how a company ends up fighting its own people.**

Leaders cannot explain the technology

People are being asked to trust decisions their leaders cannot interrogate themselves. Employees notice.

Technical fluency does not mean every executive learns to code. It means leaders understand enough to ask good questions, make tradeoffs, and explain why a decision was made.

70% of the variance in team-level engagement traces to the manager. A strategy your people never hear their own manager explain arrives as uncertainty.

Gallup

Costs you: trust, velocity, and good judgment.

Talent is organized around jobs, not capability

AI changes the boundary between what humans do and what machines should do. That should create room for people to spend more time doing what they are unusually good at. Instead, many organizations begin with headcount reduction and wonder why adoption collapses.

18 points separate employees

whose manager actively supports their team's AI use (48% engaged) from those whose manager does not (30%). Support is the variable, not access to tools.

Gallup, 43,262 US employees, 2026

Costs you: your best people disengaging from the change you need them to lead.

Nobody agreed on the intent

The board wants growth. Finance wants efficiency. Legal wants control. Technology wants capability. Employees want to understand what happens to their work. All are reasonable. Without a shared intent, every group optimizes a different answer and calls the disagreement an adoption problem.

15 points is the engagement

advantage held by organizations whose people say there is a clear plan for integrating AI into their work. Agreement is measurable.

Gallup, 43,262 US employees, 2026

Costs you: consensus by escalation, and innovation that disappears under pressure.

Weeks 1–3

Read the organization. Map the decisions, workflows, incentives, informal authority, technical constraints, and the places where people already route around the official process. Part architecture review, part organizational anthropology.

You get: an AI Intent Charter, a decision map, and an honest view of where alignment is breaking.

Weeks 4–10

Build agreement. Put leaders, technologists, operators, risk owners, and the people doing the work around the same decisions. Define what AI is for. What it is not for. Who owns the result. Where humans stay in control. What happens when someone disagrees. Leaders build enough technical understanding to lead the change rather than simply sponsor it.

You get: clear decision rights, working escalation paths, and a leadership team that explains the same strategy.

Weeks 11–16

Make it survive contact with reality. Move the decisions into the operating rhythms the organization already uses. Test what happens when the model is wrong, the employee disagrees, the regulator asks, the executive leaves, or the market changes. Then transfer the work to your people.

You get: an organization that can keep innovating when something goes wrong, not one that needs another transformation program.

WHAT YOU OWN AT THE END

A shared intent

Your people can say what AI is supposed to accomplish, where the boundaries are, and why. Not identical opinions. Shared direction.

Leaders who can lead technical change

Enough fluency to decide without hiding behind the CIO, the vendor, or the model. Enough empathy to understand what people need in order to move.

Innovation that survives

The goal is not a successful pilot. It is an organization capable of making a new thing normal, and retaining the knowledge, standards, and collective intelligence that made it possible.

WHY I DO THIS DIFFERENTLY

I started my career in ruins. Archaeology turned out to be useful preparation for enterprise technology.

Organizations leave artifacts everywhere: systems, incentives, language, workarounds, reporting structures, and the things everybody knows but nobody puts in the PowerPoint.

Later, as a Distinguished Engineer, I spent years accountable for technical outcomes without the luxury of simply managing everyone required to produce them. I had to learn how to move an organization through technical understanding, evidence, empathy, and agreement. That is the muscle I bring to this work.

I keep the advisory to **three seats at a time** because it depends on direct access to the people making the decisions. If I cannot stay close enough to the work to change the outcome, I should not be on the retainer.

The diagnostic scope is decided on the first call. I keep three advisory relationships at a time so the work stays close to the leaders and teams making the change.

This is a fit for a live adoption problem with an accountable sponsor. It is not a fit for generic AI training, a communications campaign, or transformation theatre.