

烈火

A FIELD GUIDE FOR OPERATIONS LEADERS

A I R E K K A

How to tell where AI **actually** belongs in your operation.

Before you can decide where AI fits, you have to see how your work really moves. That map is worth having on its own — whatever the answer about AI turns out to be.

烈火 THE REAL PROBLEM

The market is selling you AI. Your operation still has to tell you if it belongs.

Boards want an AI story. Vendors want a signed contract. Meanwhile your team is still reconciling three systems every Monday, hunting for a document someone swears they filed, and re-keying the same data into a portal by hand. That gap — between the AI everyone is talking about and the work as it actually runs — is where budgets quietly die.

Here is the part most people miss: **the first valuable thing isn't AI at all. It's finally seeing how your own work moves.** Most organizations have never had their real process written down — not the tidy diagram from the SOP binder, but the true path, including the unofficial steps, the workarounds, and the one person everyone quietly depends on. Put that map on the table and decisions get easier immediately, whether or not a single thing ever gets automated.

That map is the point of an audit. AI is what we evaluate *against* it — honestly, including the times the answer is "leave this alone."

流れ WHY MOST "AI ASSESSMENTS" ARE THEATER

The discipline underneath the buzzwords is older and simpler than the branding.

There is a lot of mystique around AI assessments — frameworks, certifications, colored belts, proprietary-sounding acronyms. Strip the branding away and the real work is something that has underpinned every serious systems project for forty years: sit with the people who do the job, map how the work actually moves, find where it breaks, and judge honestly what is worth changing.

That discipline is called **systems analysis**. The industry keeps renaming its pieces — value stream mapping, process discovery, requirements elicitation, feasibility scoring — but the fundamentals don't change. We run those fundamentals. Where a standard name helps us talk to your team, or where a regulated environment needs something people recognize, we use it. But the value was never in the acronym. It's in someone looking clearly at your operation and telling you the truth about it.

Here is the lens we use. You can start applying it to your own operation right now.

六 SIX LENSES ON THE WORK

01 • WORKFLOWS

The real path of work

Where does time actually die? We map the true route a task takes — including the steps that live in nobody's documentation — and look for where exceptions pile up and where two teams argue about whose number is right.

Usually called: current-state process analysis · value stream mapping · process discovery

ASK YOURSELF

If you followed one important task from start to finish, how many handoffs would surprise you — and how many happen in an inbox or a spreadsheet nobody planned for?

02 • SYSTEMS

What you already run — and where a person is the integration

We inventory the systems you have and, more importantly, how they connect. Often the "integration" between two systems turns out to be a human quietly re-typing data from one into the other.

Usually called: systems and integration analysis · application portfolio assessment

ASK YOURSELF

Where in your operation is a person the bridge between two systems that were supposed to talk to each other?

03 • DOCUMENTS

What's trapped in files

PDFs, scans, emails, portals, records. We trace the unstructured material into the actions it feeds. If people are reading documents so a system can do something, that is exactly where the friction — and the opportunity — hides.

Usually called: data and information flow analysis

ASK YOURSELF

How much of your team's day is spent reading something in one place so they can enter it in another?

04 • PEOPLE

Who handles the messy cases

We talk to the people doing the work and find out what they actually decide — especially on the exceptions. Some of that judgment can be supported by a system. Some of it must stay human. Knowing the difference is most of the job.

Usually called: requirements elicitation and exception analysis · stakeholder discovery

ASK YOURSELF

Which decisions in your operation live entirely in one experienced person's head — and what happens the week they're out?

05 • RISK

Whether it can be trusted here

We look at data quality, integration cost, irreversible actions, and compliance exposure — not whether a demo looks impressive, but whether a system can be trusted in *your* operation, at *your* tolerance for error. Sometimes the right answer is keeping data in-house on a local model rather than a public tool.

Usually called: feasibility and risk assessment · aligned, where it matters, to the NIST AI Risk Management Framework

ASK YOURSELF

If a system got one of these decisions wrong, could you undo it — and would you know it had happened?

06 • SEQUENCE

What to do first

Findings are only useful in order. We turn them into a sequence: what to do first, what to fix before AI touches it, and what to leave alone. Not a shopping list of tools — a ranked plan you can actually follow.

Usually called: gap analysis and prioritized roadmapping

ASK YOURSELF

If you had to pick the one process whose fix would free up the most time, could you defend the choice with evidence — or is it a guess?

実証 THE THREE VERDICTS

Every candidate earns one of three honest calls.

● Go

The process is stable enough, the data is reachable, and the payoff is real. This is a candidate for a scoped pilot — built carefully, with human gates and an audit trail.

● Fix first

AI would only amplify a mess. Clean the handoff, the master data, or the exception path first — then re-score. The map tells you exactly what to fix.

● Don't automate

Low volume, high judgment, thin data, or unacceptable risk. Leave it with your people. That is a finding, not a failure — and knowing it saves you from paying to automate something that shouldn't be.

例 WHAT THIS LOOKS LIKE IN PRACTICE

The mailbox was the ticketing system. The extra software was just another hop.

A common pattern: support already lives in a shared mailbox. Customers write in, and staff either work the inbox directly or copy each message into a separate ticketing tool — running two desks to do one job. The mailbox was the real system all along; the extra tool was the detour.

Mapping the work first changes the recommendation entirely. Instead of bolting a chatbot onto the ticketing layer and hoping the detour disappears, the honest answer is: **start where the work already arrives**. Pull messages at the mailbox, let a well-bounded system resolve the routine ones, and *escalate everything else to a person*. That gate — close what can be closed, hand the rest to a human — is the whole point.

You only see that clearly once the work is mapped. Lead with a tool, and you'd have automated the detour instead of removing it.

烈火 WHAT IT'S WORTH — AND WHAT COMES NEXT

The map is the deliverable. It's yours to keep.

An audit ends with a written brief you own: how your work actually moves, where AI belongs, where it doesn't, and what to do first. Even if the verdict on every item were "don't automate," you'd still walk away with something most organizations have never had — an honest, current picture of their own operation. That clarity is the value. The AI recommendations sit on top of it.

You can act on that brief however you like: run it internally, hand it to someone else, or ask us to build the first thing that earned a go.

If something earns a "go," the build is its own study.

An audit tells you **whether** and **what**. Building the solution requires getting closer — because integration reality only reveals itself first-hand. How your systems truly pass data, where the edge cases hide, what breaks under real load: those are answered by inspecting the work directly, not from the map alone. That's why a build is a separate, more thorough engagement with its own scope and its own quote. It's also how a solution ends up fitting your operation like a glove instead of like a template.

Start with a conversation — free, no pitch. Tell us the process eating your week at airekka.com/contact

AIREKKA performs operational AI audits, then recommends where AI makes sense — and where it does not. Manufacturing, healthcare, insurance, and other operations where disconnected systems are costing real money. U.S.-based · Toledo, OH & Orlando, FL · airekka.com