PREPARED FOR THE AI SYMPOSIUM FOR BUSINESS AVIATION

Selecting Your AI Partner

A field guide for operators choosing who to build with.

INSIDE THIS PLAYBOOK

01 Why This Industry Is Different

Generic AI advice assumes a software company. You are not one, and most of the differences work in your favour once you plan around them.

A safety culture already exists
You run checklists, briefings, and post-event reviews. That discipline is exactly what makes a small pilot succeed, because you already know how to scope a change and watch it.

Real regulatory exposure
Part 91, 121, 135, 145 and 141 all sit somewhere in this room. Some decisions never get automated, and a good partner will say so before you ask.

Data lives in many places
Scheduling, maintenance tracking, fuel, accounting, and several inboxes. Very little of it connects, and nobody knows its true shape until somebody looks.

Small technology teams
Often one person, sometimes none, and usually already busy. Any plan assuming spare internal engineering capacity is a plan that will not happen.

Peaks, not averages
The work does not arrive evenly. Hours recovered in a quiet week are worth considerably more in the week when everything happens at once.

Relationships are the product
Charter, maintenance, training, ground handling. In all of them, the person who remembers a customer by name is doing something no software replaces.

THE MISTAKE THAT COSTS THE MOST

Buying the end state before building the habit.

The organisations that do well here start with ordinary, low-risk work and a handful of people using AI properly on real tasks. The ones that struggle buy an ambitious system first, then discover nobody has the time, the data, or the confidence to operate it. **Readiness is not a prerequisite you can purchase.**

> WHAT TO TAKE FROM THIS
>
> None of these is a reason to wait. They are reasons to sequence properly: small scope, a named owner, and a person who checks the output before it matters. Judge a provider on whether they will meet you where you are today rather than where their product assumes you already are.

WHERE TO BEGIN, PRACTICALLY

Answer the next page honestly before anyone answers it for you.

The readiness scorecard opposite is fifteen boxes you can tick in about ten minutes with two or three colleagues. It is deliberately about you rather than about any product, because the gaps it surfaces are the ones a provider cannot close on your behalf. **Every unchecked box is either work to do first, or a question to put to whoever you are about to hire.**

AI Readiness Scorecard

Check every box your organisation can honestly claim today. One point per item. Nothing here requires buying anything, and every unchecked box is one no partner can fill for you.

TOTAL SCORE ____ / 15

PEOPLE

AI succeeds or fails on the people behind it. Leaders set the tone for everyone else.

- [ ] **Executive Sponsorship** — Leadership is actively sponsoring AI, not merely approving it, so teams feel safe to explore and adopt.
- [ ] **Upskilled Leaders** — Senior leaders understand AI well enough to champion responsible use and make informed decisions.
- [ ] **Informal Champions** — Early adopters inside your functions are identified and activated as peer advocates.
- [ ] **Role-Based Learning** — Training is built around what each role actually does, not generic AI literacy.
- [ ] **Psychological Safety** — People feel safe raising concerns, asking basic questions, and admitting mistakes.

PROCESS

Change management turns AI ambition into adoption. It is where most efforts quietly succeed or fail.

- [ ] **Clear Vision** — There is a shared story of where the organisation is going with AI, and why it matters.
- [ ] **Stakeholder Alignment** — Leaders, teams, and influencers are brought in early, so resistance becomes momentum.
- [ ] **Structured Rollout** — AI is introduced in deliberate waves so teams can absorb it and build confidence.
- [ ] **Human in the Loop** — Review gates are built into every process so judgment and accountability stay with people.
- [ ] **Improvement Loop** — A feedback mechanism captures what is working, so processes improve rather than stall.

TECHNOLOGY

The technical foundation differs from traditional software. Getting it right makes the rest possible.

- [ ] **Data Connectivity** — Core systems are connected so AI can work across the operation without manual exports.
- [ ] **Trustworthy Data** — The data powering AI is accurate and consistent enough to act on with confidence.
- [ ] **IT Partnership** — People and technology teams agree on priorities, so work does not stall waiting on capacity.
- [ ] **Security and Governance** — Data is protected with proper access controls, and AI use meets privacy and legal requirements.
- [ ] **Oversight Process** — A defined process reviews what the technology produces, with clear ownership for flagging errors.

0 – 2
ALIGNED
AI is on the radar and the conversation has started, but foundational work has not.

3 – 5
PREPARED
Early groundwork is in place and key leaders are engaged, but gaps remain across most pillars.

6 – 8
ADOPTED
Tools are in use and teams are building confidence, but nothing has scaled consistently.

9 – 11
SCALED
AI is embedded across functions with governance and measurable impact. Isolated gaps remain.

12 – 15
EVOLVED
AI is core to how the organisation operates and improves. You are positioned to lead.

What You Can Do Without Hiring Anyone

Most technology arrives as a product: you buy the capability. AI is unusual in that a large share of the capability arrives with trained people using tools your company may already pay for. That changes what you should buy, and when, and it is the most useful thing to settle before you take a single vendor meeting.

DO THIS YOURSELVES — WEEKS, NOT QUARTERS

Drafting and editing
Proposals, emails, reports, policies, job descriptions. A first draft in minutes, edited by the person who owns the outcome.

Summarising and synthesis
Long threads, transcripts, reports, and regulatory notices turned into something a busy person will actually read.

Research and first-pass analysis
Market questions, supplier checks, comparing options. Done before you commit anyone's week to it.

Questions of your own documents
Contracts, manuals, policies, records. Upload and ask. No integration and no project plan required.

THIS NEEDS HELP — WHERE A PARTNER EARNS THE FEE

Connecting systems
Anything that reads or writes across your scheduling, accounting, or operational software rather than sitting beside it.

Anything that acts
Agents that do rather than draft, and the accountability structure that has to exist around them before they do.

Customer data at scale
Privacy, access control, and the governance that keeps a decision defensible six months after it was made.

Oversight design
Who reviews what, how errors surface, what gets logged, and who is named when something goes wrong.

THE HONEST MIDDLE

There is a third category, and nobody puts it on a slide.

Plenty of work you could do yourselves, you will not, because nobody has the time to own it. That is a legitimate reason to buy help, and a far better one than believing the work is technically beyond you. **Be clear which you are in: the first is a capacity problem and the second is a capability problem, and they call for very different conversations.**

> **The best first money is usually training, not software. Spend it on the people who will live with the result.**

SECTION TWO

Selecting Your Partner

Almost every provider is strong in one of three areas and quieter about the other two.

02 Selecting Your Partner

Most providers you meet will be genuinely good at something. The useful question is not whether they are capable. It is which of these three you are talking to, because each is the right answer to a different problem.

The Niche Provider

Deep domain knowledge, limited engineering.

Knows your business well, often already sits in your stack, and is fast and inexpensive to start with. The AI is usually a capability added to a product you already own. **Where it stops:** it works on the data inside their product and little else, and anything outside that lane becomes a roadmap date.

The Development Shop

Real engineering, less industry context.

They can build what you describe, on a schedule, with real engineers. Genuinely valuable when you already know precisely what you need. **Where it stops:** they build exactly what was specified, and nobody in the room is positioned to tell you the specification is wrong.

The Consultant

Strong judgment, thinner delivery.

Excellent diagnosis and real strategic value in the first six weeks. You get a roadmap, a maturity model, and a clear view of the opportunity. **Where it stops:** the deliverable is a document, and documents do not reduce anyone's workload on their own.

ONE QUESTION THAT SORTS THEM QUICKLY

Ask what they would build if you gave them a month and no instructions. The Niche Provider describes their existing product. The Development Shop asks you for requirements. The Consultant proposes a discovery phase. The partner you want describes a specific workflow inside your business and explains why it is costing you hours.

WHAT THEY HAVE IN COMMON

All three are the right buy for somebody.

Each is genuinely excellent at the thing it was built for, and for a narrow enough problem any of them is the correct choice. **The mistake is not picking the wrong archetype. It is paying for one and expecting another,** which is how a perfectly good provider ends up in a difficult relationship eight months later.

What You Are Actually Looking For

The partner you want sits in the overlap: someone who understands how your business makes money, can build the thing, and will tell you when the thing is a bad idea. Two of the three is the most common shape in this market, and it is workable once you know which one is missing and who is covering it.

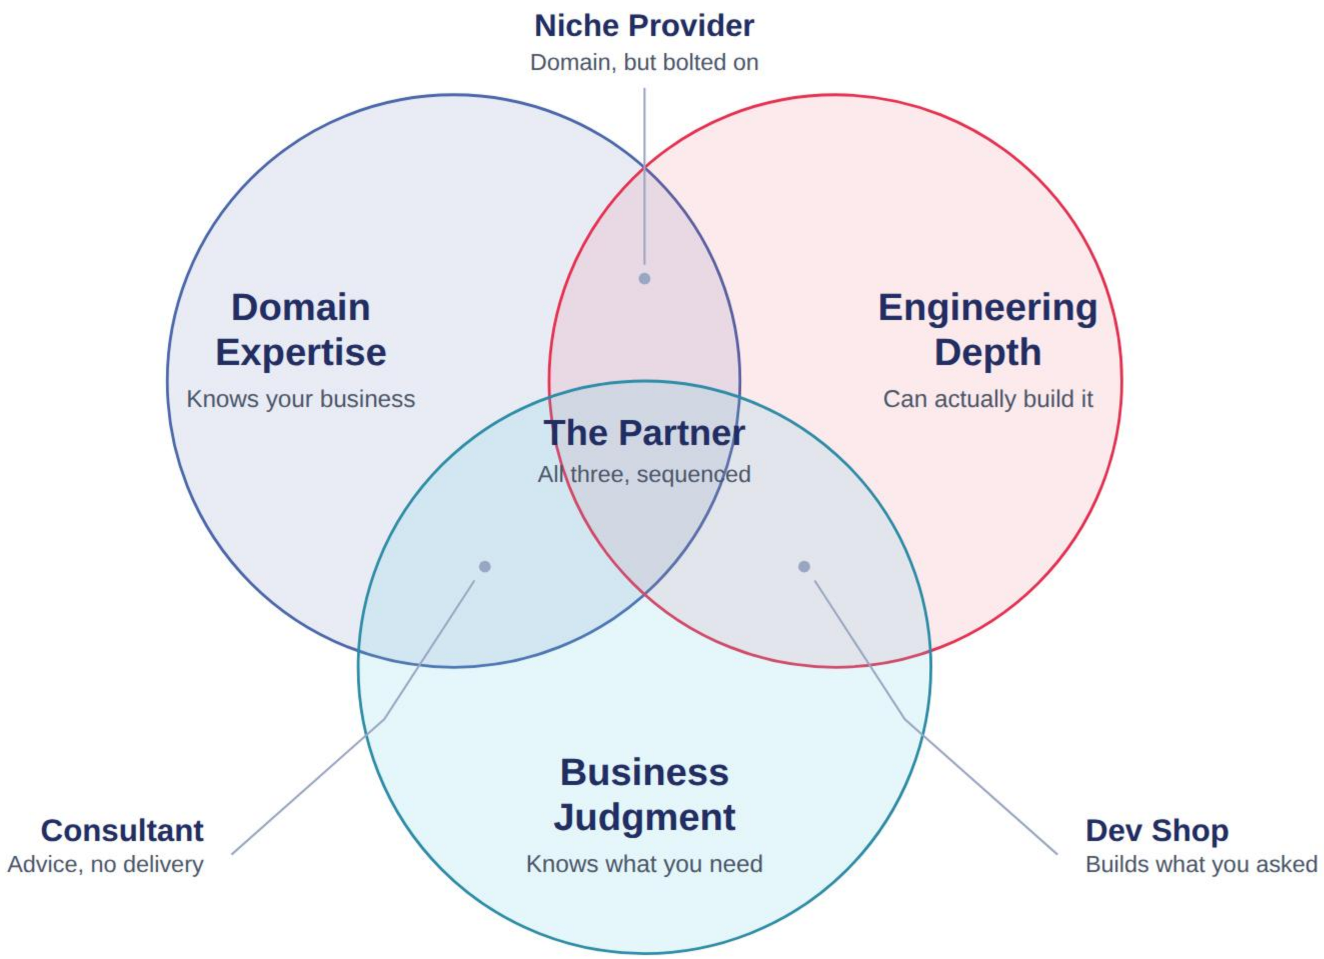

A FAIR QUESTION TO ASK OUT LOUD

Ask which two of the three they have, and which one they partner for. It is not a trap, and most good providers answer it readily because they have already had the conversation internally. The answer tells you what you will need to supply yourself, which is far more useful than a claim that they cover everything.

Reading the Market

The AI market is moving faster than anyone's ability to evaluate it. Very little of what follows is bad faith. It is a young industry describing its best case to buyers who have no established way to check it. This playbook is a way to check it.

WHAT GETS SAID	WHAT USUALLY HAPPENS
"AI will run your operation"	AI drafts. People decide. Every regulated call stays human.
"Plug it in and it learns your business"	It learns nothing until someone curates what it reads.
"One platform for everything"	Two or three narrow tools, run well, beats one run badly.
"Replace headcount"	Recover hours. The people you have get their evenings back.
"This is a six figure transformation"	The first real win is usually small, ordinary, and fast.
"It makes things up, so it is not safe here"	Fair, and still incomplete. Scoped and supervised, the risk is manageable. The version nobody owns is not.
"AI is a project"	It is an operating habit. Projects end. Habits compound.
"Our model is 99 percent accurate"	On whose data, measured how, and who finds out about the rest?
"Your team will love it"	Your team will use it if it removes work they dislike. Otherwise it is another tab nobody opens.

HOW TO USE THIS PAGE

Take it into your next provider conversation and listen for the left column. Nobody says these sentences quite this bluntly, but the claims underneath are the same, and hearing three of them in one meeting tells you what kind of conversation you are in. None of them makes a provider wrong. They mean the real answer has not arrived yet.

The useful question is not what the software can do. It is what it can do here.

The Question Bank

Eleven questions, each with what a good answer sounds like, because the question is worth little if you cannot tell a real answer from a confident one. None of these is hostile, and a strong provider will enjoy most of them.

The proof question

Show us something in production at a business our size.

A good answer: A named customer, a specific workflow, a number attached to it, and an offer to connect you.

A weak answer: A highlight reel, a logo wall, or confidentiality cited with no anonymised alternative.

The boundary question

What would you decline to automate for us?

A good answer: Immediate and specific, usually about safety, judgment, or the moment a customer is upset.

A weak answer: Hesitation, or a version of "it depends on your risk tolerance."

The ownership question

Who owns the data, the prompts, and the outputs if we part ways?

A good answer: A direct answer, in writing, that survives the conversation and appears in the agreement.

A weak answer: "That is handled in the contract," without producing the clause.

The training question

Is our data training your model, or anyone else's?

A good answer: No, with a subprocessor list to back it up.

A weak answer: "Anonymised" or "aggregated" without explaining what those words mean in their contract.

The team question

Who will we actually work with, and what do they know about our industry?

A good answer: Names the delivery team and offers to bring them to the next call.

A weak answer: "You will be assigned a dedicated success manager."

HOW TO USE THESE

Ask them plainly, early, and in this order. The first five establish whether there is anything real underneath, and they are fair questions that a confident provider will answer without hesitation. Write the answers down in the meeting; by the third conversation the details blur and only your notes will still be accurate.

The Question Bank, Continued

The second five are about what happens after the agreement is signed, which is where most of the real difference between providers shows up.

The scope question

What is the smallest useful thing you could deliver in thirty days?

A good answer: A narrow, ordinary, specific workflow with a named owner on both sides.

A weak answer: A platform, a pilot programme, or a phase one that ends in a readout.

The endurance question

What happens in month seven, once the novelty has worn off?

A good answer: They have seen it, they name the failure pattern, and they have a plan for it.

A weak answer: They have never been in a month seven.

The integration question

How deep does this go into our systems, and what happens if you disappear?

A good answer: An honest map of the integration and a clear account of what you would still own.

A weak answer: Vagueness about dependencies, or an architecture where leaving means starting again.

The accountability question

When it gets something wrong, who finds out, and who owns it?

A good answer: A named person, a review step, and logging you can audit.

A weak answer: “The model is highly accurate.”

The scale question

What does this cost when it works, and when it works ten times more?

A good answer: Both numbers, offered without being asked.

A weak answer: The first number only.

The independence question

What will our team be able to do without you in six months?

A good answer: Specific about what they will teach you, what they hand over, and what they expect to still be doing.

A weak answer: Treats it as an odd question, or answers only with what their platform will be doing by then.

If a provider cannot answer these in a first meeting, that is information rather than a verdict.

SECTION THREE

Tools You Can Use

Two pages you can print, fill in, and take into the room with you.

03 AI Partner Scorecard

The same fifteen boxes, in the same three pillars, pointed outward. Score each provider in the same week, before the conversations blur together.

TOTAL SCORE ____ / 15

PEOPLE

You are buying a team, not a product. These five tell you who you will actually work with.

- ☐ **Named Delivery Team** — The people who will do the work are named and available, not only the people in the pitch.
- ☐ **Industry Understanding** — Someone on their side understands how a business like yours actually makes money.
- ☐ **Accessible Escalation** — You know who to call when something is wrong, by name, rather than a shared queue.
- ☐ **Honest Self-Assessment** — They can say what they are not good at without being pushed into it.
- ☐ **Leaves You More Capable** — They teach your team as they build, so you are not routing every small change through them forever.

PROCESS

How they work matters more than what they demo. These five predict month seven.

- ☐ **Proven in Production** — They can show working software at a comparable business, not a demonstration environment.
- ☐ **Small First Delivery** — They proposed something genuinely useful they could ship in about thirty days.
- ☐ **Clear Boundaries** — They named something they would decline to build for you, and explained why.
- ☐ **Error Accountability** — When it produces a wrong answer, a named person finds out and owns the fix.
- ☐ **Life After Launch** — They have a plan for the months after go-live, and have lived through them before.

TECHNOLOGY

The terms and technical posture behind the work. These five are the ones that bite later.

- ☐ **Data Ownership** — Your data, prompts, agents, and outputs are yours, stated in writing.
- ☐ **No Training on Your Data** — Explicit confirmation your data trains neither their model nor anyone else's.
- ☐ **Integration Honesty** — They asked about your systems and your data before quoting a price.
- ☐ **Security and Access** — Where data lives, who can see it, and what happens in a breach, all answered plainly.
- ☐ **A Real Exit** — You can leave with your data and your workflow logic intact, on a known timeline.

0 – 2
NOT A FIT
Very little to evaluate yet. Keep the relationship, and keep looking.

3 – 5
UNPROVEN
Possibly capable, but nothing here you could defend to your board.

6 – 8
PROMISING
A genuine candidate. Take a second meeting and push on every unchecked box.

9 – 11
STRONG
A serious partner. Negotiate the remaining gaps into the agreement.

12 – 15
RARE
Unusual in this market. If the references hold up, move.

Contract Terms Checklist

Ten terms worth settling before signature. Every one is something operators most often discover was missing only after something went wrong, and all of them are far cheaper to agree now than later.

Data ownership
Your data is yours during the engagement and after it, in a usable export format rather than a report or a screenshot.

Model training rights
An explicit prohibition on training their models, or any third-party model, on your data.

Subprocessors
A named list of everyone who touches your data, and notification before that list changes.

IP in what gets built
Who owns the prompts, the agents, the configurations, and the workflow logic. Most agreements are silent here.

Exit terms
Notice period, transition assistance, export format, and what happens to your data on their side afterwards.

Support definition
What support means in hours and response times, in writing. Best efforts is not a support model.

Scale pricing
The price per unit at ten times your current volume, locked or capped, so success does not break the budget.

Security posture
Where the data physically lives, who inside their company can see it, and what happens in a breach.

Insurance
Appropriate to the level of access you are granting them.

Renewal terms
No multi-year automatic renewal before anything has actually shipped.

THE ONE THAT MATTERS MOST

If the agreement has no exit path that leaves you with your own data and your own workflow logic, there is little else worth negotiating. Everything else on this page is a term. That one is the difference between a partner and a dependency.

WHO SHOULD SEE THIS FIRST

Ten terms, one afternoon with your attorney.

Most operators review an AI agreement on the provider's paper, which means every silence in the document is a silence somebody else chose. Arriving with your own list changes what gets discussed. **Considerably cheaper now than in the month you decide to leave.**

SECTION FOUR

Starting Well

The first win is almost always small, ordinary, and fast.

04 Starting Well

The filter matters more than the list. Start where the work is repetitive, high in volume, and low in consequence if a person checks the output. Almost every function in your business has work that fits, and the first win arrives fast enough to keep everyone interested.

START HERE — LOW RISK, HIGH VOLUME, EASY TO MEASURE

Finance and accounting
Receipt capture and coding, invoice matching, reconciliation, and the month end assembly that somebody currently retypes between systems.

Recurring reporting
The weekly and monthly packs rebuilt by hand every cycle. Assembled once from the source, then reviewed by the person who signs it.

Finding things
Contracts, manuals, policies, records, past proposals. Ask a question in plain language instead of knowing exactly where to look.

Meetings and follow-through
Notes, summaries, and action items captured and routed, so the commitments made in a room survive the week after it.

THEN HERE — HIGHER VALUE, ONCE THE HABIT AND THE REVIEW STEP HOLD

Sales and business development
Proposal and quote drafting, research before a call, follow-up that would otherwise go quiet, and keeping the CRM honest without anybody policing it.

Operations and scheduling
Inbound request triage, conflict detection, exception flagging, and the dozen small coordination tasks that fill a coordinator's day.

Marketing and communications
First drafts, repurposing one piece of work across channels, and customer communications that a person edits rather than writes from nothing.

Executive intelligence
Pulling a picture together across finance, operations, and sales for a decision, in an afternoon rather than a fortnight of asking people for numbers.

WHY THE ORDINARY WORK GOES FIRST

You are buying judgment, not output.

The first weeks are how your people learn where the tool is strong and where it is confidently wrong, and it is far cheaper to learn that on an expense report than on a customer. By the time you reach the second list, somebody in the building can look at an output and tell within seconds whether it is right. **That person is the actual asset, and no provider can sell you one.**

Small, ordinary, and fast beats ambitious and stalled. The second list is not harder technology, it is higher consequence.

What Stays Human

Everything on this page is something you could automate today. That is the point. Just because you can does not mean you should, and in a business built on relationships this list runs well beyond what the regulations require.

The yes is human
A significant purchase is a relationship decision. Draft the quote all day long. Never draft the yes.

Delivering bad news
A delay, an overrun, something you got wrong. A name goes on that call, and an apology from software is worse than none.

Setting the standard
Culture is what you tolerate, and no system tolerates anything. The standard is set by what a person declines to walk past.

The awkward ask
The favour outside the normal process. The answer is a judgment call, not a configuration setting.

The unpopular call
Somebody has to own the decision nobody wanted, then stand in front of the people it affects and explain it.

Nobody reports to software
Your team follows a person they respect, not a task queue. Route work through systems; do not manage people with them.

Hard conversations
Pay, performance, and difficult feedback. A person, in a room, every time.

Saying no
Declining work, a request, or an idea. Your name on it, said out loud, is the whole job.

Recognition that lands
Thanks from a dashboard is not thanks. People can tell the difference instantly, and they remember which one they got.

Noticing something is off
The experienced person who senses a problem before any system flags it. Never build a process that overrules that instinct.

THE TEST TO APPLY

Ask who is on the other end of it, and how they would feel if they learned a system had handled it. If the honest answer is that a customer or an employee would feel diminished, you have found your line. It is a better test than any rule, because it travels to situations this page has not thought of.

WRITE THIS LIST DOWN BEFORE YOU BUY ANYTHING

The exercise takes an hour and it changes every conversation that follows, because you are no longer being shown what is possible, you are asking what belongs inside a boundary you have already set. A provider who will not draw this line with you is telling you something worth hearing.

Automate the work around the relationship. Automate the relationship itself and you have paid to become worse at it.

Your First Ninety Days

The shape matters more than the calendar. You are deliberately climbing a risk curve: ordinary work first so people build judgment cheaply, then one measured workflow, then the things that actually move the business, once the habit and the review step can carry them.

Train first · low risk work

Higher value · once it holds

1

DAYS 1–30

Train your people

Eight or ten people using AI properly on ordinary tasks, in tools you already pay for. Nothing customer facing or regulated. Ground rules first.

2

DAYS 31–60

Prove one workflow

Measure it before you change it: hours, volume, errors. Build the smallest useful version and keep the manual process running beside it.

3

DAYS 61–75

Make it how you work

Write down the review step: who checks output, how often, and what happens when it is wrong. Teach a second team the same pattern.

4

DAYS 76–90

Move up the curve

Take on something with real weight. Decide on the first build from evidence: expand, stop, or move on. Choose the next two with the scorecard.

THE STEP ALMOST EVERYONE SKIPS

Measure it before you change it.

Two hours spent counting hours, volume, and error rate in week five is what turns day ninety into a decision rather than an argument about who feels most positive. **A provider who lets you skip the baseline is not saving you work. They are removing the evidence.**

WHAT GOOD LOOKS LIKE AT DAY NINETY

One workflow measurably better, a documented review step, a second team running the same pattern, and enough judgment in the building that somebody can look at an output and know within seconds whether it is right. Not a transformation. A foundation that the next twelve months can stand on.

Ninety days does not make you an AI company. It makes you an organisation that knows how it will decide.

PREPARED FOR AISBA 2026

Make work
more meaningful.

This playbook was written for the AI Symposium for Business Aviation and is free to use, copy, and share inside your company. No form, no email capture, no strings.

Prepared by Jason Rhoades, Meaningful AI · jason@meaningful.ai · meaningful.ai