



ADVANCED TRACK · SESSION RECAP

AI Symposium for Business Aviation (AISBA)

What the advanced track covered, and what to do with it next.



INSIDE THIS RECAP

- 01 The Room, and the Gap**
Where adoption actually sits, and the 80/20 inside every role.
- 02 The Three Gears**
How capability climbs, and why leadership goes first.
- 03 Inside the Live Build**
Four agents, live systems, thirty minutes.
- 04 Authority, Not Enthusiasm**
Governance that keeps pace with capability.
- 05 What to Do Next**
The three moves worth making this quarter.

Who Presented

ABOUT THIS RECAP

The advanced track ran ninety minutes and deliberately mixed three things: a framework for where an organization actually sits, a live build against a real charter operation's systems, and the governance discipline that has to keep pace with both. This document is the durable version of that session, written so it can be passed to someone who was not in the room.



Jason Rhoades

YOUR AI ADVANCEMENT

Opened the session with the adoption gradient, the 80/20 inside every role, and the three gears.

jason@meaningful.ai



Diego Magrini

BUILDING AI AGENTS

Built four working agents live against a real Part 135 operation, start to finish, in front of the room.

diego@nerd.aero



Jiri Marousek

SAFETY, SECURITY AND COMPLIANCE

Grounded the session in aviation safety, security and compliance systems, and where AI meets them.

jiri.marousek@airtera.com



Toby Benenson, AIGP

GOVERNANCE AND AUTHORITY

Founder of SayFlight. Closed with the participation taxonomy, the acceptance test, and what operational authority requires.

toby@sayflight.ai

Three perspectives, one argument: capability is already here, and **authority is what has to catch up.**


01 The Room, and the Gap

After two years of speaking on this subject, the hard part is no longer explaining the technology. It is the spread. Experience, perception and access widen every day, and a single room now holds people who have never opened a frontier model alongside people running agents against live operational systems.

INSIDE ANY ROLE 80% Clerical work that can be handed to an agent.	WHAT REMAINS 20% Relational work, and it rises in value.	AVAILABLE TODAY 1 day Per week, recoverable with current tools.
---	--	---

Where the world actually sits

71% HAVE NEVER USED IT of humanity has never knowingly used generative AI. <i>DataReportal, April 2026</i>	17.8% USE IT AT ALL of the world's working-age population, up 1.5 points in a single quarter. <i>Microsoft AI Diffusion Report, Q1 2026</i>	1% HAVE ACTUALLY FINISHED of executives call their rollouts mature, though 88% have deployed something. <i>McKinsey, 2025</i>
--	---	---



Microsoft AI Diffusion Report, first quarter 2026. Twenty-six economies now exceed 30 percent.

THE DISTANCE IS A DECISION

The top of that table is a day away, not a career away.

Nobody in that room needed a new degree to move from paid use into building. They needed a day of intent and a reason. The gap between the 17.8 percent who use AI and the 1 percent who have finished the job is **curiosity plus permission**, and both are free.

THE QUESTION THAT RAN THROUGH THE SYMPOSIUM

Job displacement was the loudest topic and the least useful one. The panel agreed it is overhyped. The real exposure is an education gap, and it arrives at the same moment as a retirement wave, which makes this a hiring opportunity for anyone who can build these skills rather than a threat.

The 80/20 Inside Every Role

Process mapping across operations, finance, HR, project management and leadership keeps returning the same split. It held in every role examined, at every level, with the honest exception of hands-on maintenance.



The 80 percent

Clerical, and agentifiable.

What it is: translation between systems, re-keying, cutting and pasting, chasing status, assembling the same report every month. **Why it goes first:** none of it ever required judgment, and all of it leaves a trail an agent can follow.



The 20 percent

Relational, and rising.

What it is: vendor and team management, customer relationships, the judgment calls that carry risk, the conversations that need a face. **Why it matters:** as the clerical share collapses, this becomes the whole job, and the day looks more like a whiteboard than a laptop.

WHAT ACTUALLY HAPPENS DOWNSTREAM

When one machine runs a hundred times faster, everything behind it has to keep up.

A small group of architects and engineers can now produce more than the organization around them can review. That pressure shows up as hiring, not as redundancy, and it lands on the roles that check, approve and decide.

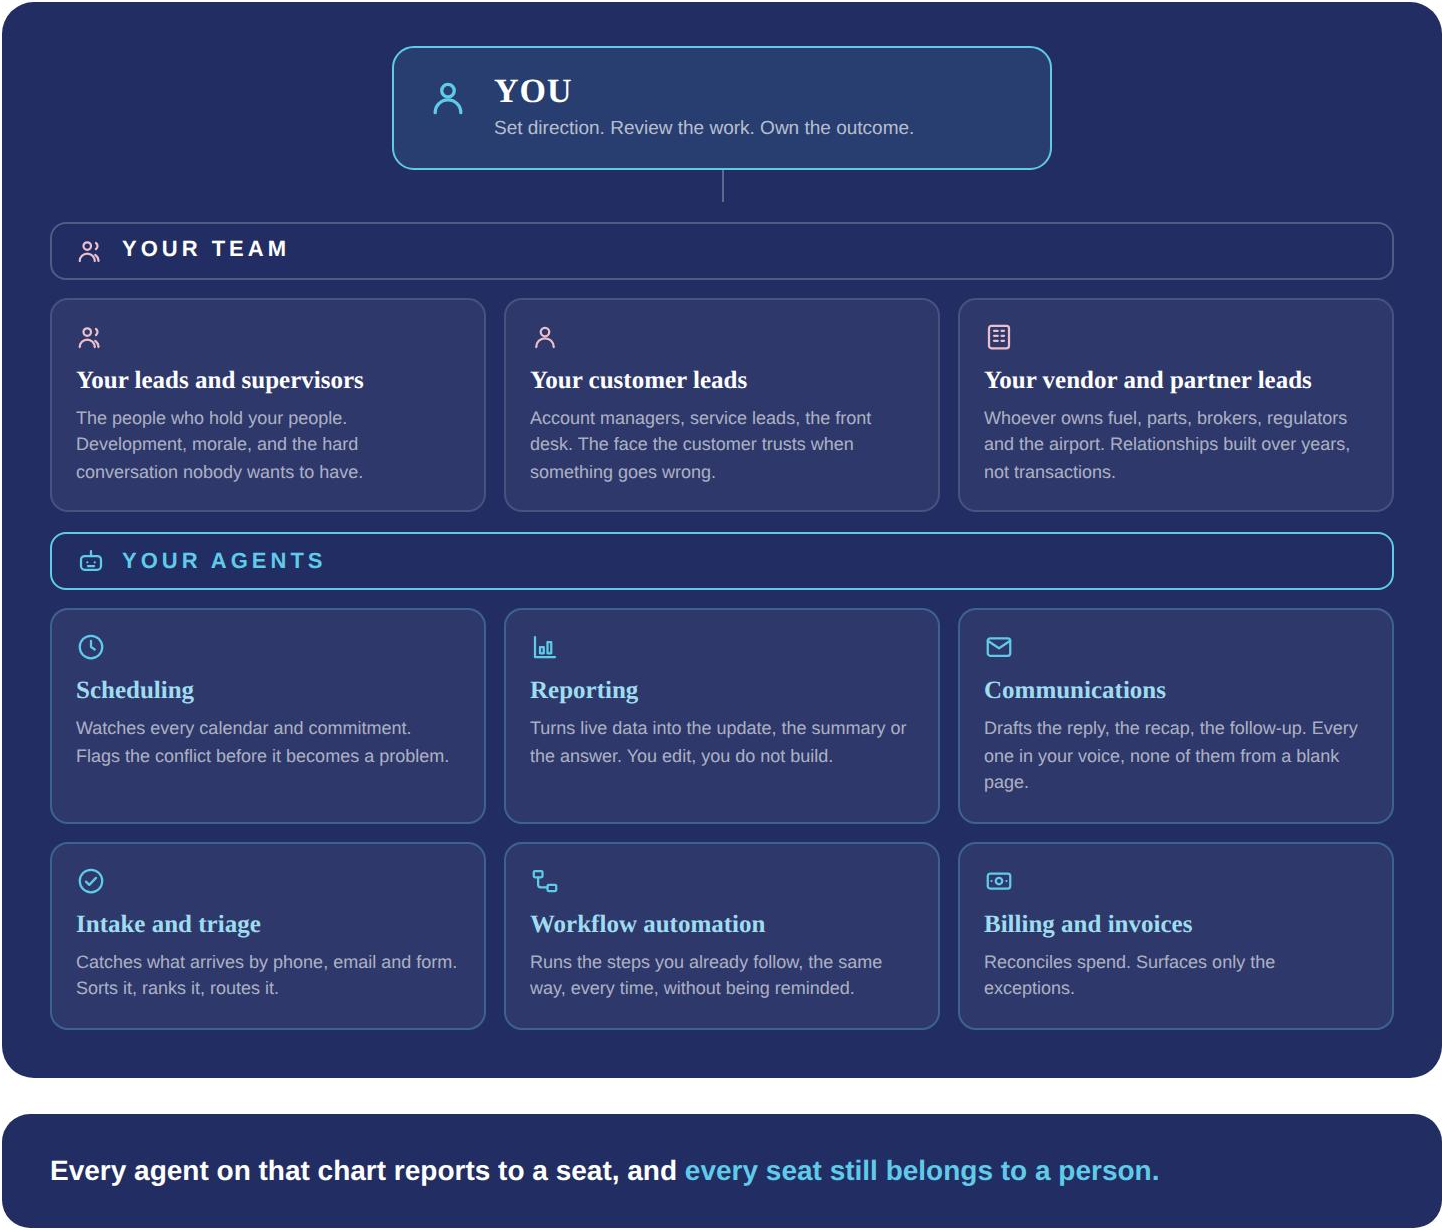
THE FINDING THAT SURPRISED US

The original assumption going into these engagements was that some roles could be replaced outright. Process mapping said otherwise. Every function examined, operations, finance, HR, project management and leadership, still needs a human attached to it. What changes is the ratio inside the role, not whether the role exists.

The leaders of the next few years will be leading humans, agents, or **very likely both.**

The Org Chart You Will Actually Run

Put the 80/20 on paper and the shape of the next few years falls out of it. A small number of people holding the relationships, and a larger bench of narrow agents behind them, each one accountable to a named human.



SECTION TWO

Capability arrives in gears, and the climb is the whole strategy.

Individual use, then shared practice, then organizational change. Most operations are sitting at one and a half.

02 The Three Gears

People, process and technology is still the flywheel, but the technology share of it keeps shrinking. Building software is converging on near zero cost and near zero time. The cost of culture is moving the other way, which is where the leadership investment belongs.



1 day

Per week, still being done by hand.

Gear one: the individual and the model

Prompting, drafting, research, summarizing. Almost everyone in a modern office is here, and an organization can coast here for a long time. Even so, the current tools already cover roughly a day a week of what people are still doing manually.

25%

More of the week spent on meaningful work.

Gear two: skills, agents and a team

Repeatable routines taught once and reused, workflows that run without being asked, and a team that exchanges them. This is where tasks fall away in volume rather than minutes, and where the practices are still genuinely emerging.

90%

Of a role, covered by systems rather than hands.

Gear three: the organization changes shape

Custom software, a private and governed data layer, decision intelligence that starts answering questions nobody thought to ask. Few operations are here. Most are at about one and a half, having dabbled without committing.

Why Gear Two Starts at the Top

Walking into an organization, the first conversation is with leadership, not with IT. That order is deliberate, and the research keeps landing in the same place: the people are ready, and the steering is what is missing.

The barrier is not the workforce

3x

MORE USE THAN LEADERS THINK
more employees use AI for a third or more of their work than their leaders estimate.
McKinsey, 2025

47%

OF THE C-SUITE
say their own companies are moving too slowly, though 69% began investing over a year ago. *McKinsey, 2025*

1%

REACH MATURITY
of leaders call their deployment mature, with steering, not staff, named as the constraint. *McKinsey, 2025*

What a Gear 2 leadership team can do

The working definition, not the aspiration

- Exchange a skill or a reusable routine with a colleague
- Share an agent, and review someone else's before it runs
- Judge whether an output is trustworthy enough to send
- Tell a real automation apart from a convincing demonstration

THE TIME IT TAKES

One quarter is enough. A leadership team that commits to Gear 2 this quarter can start Gear 3 initiatives in the next one, which is a far shorter runway than most boards assume.

The organizations that stall are not the ones without tools. They are the ones whose leaders **never learned to use them.**



SECTION THREE

Thirty minutes, four agents, live systems.

What the room watched get built, and the engineering choices that made it safe enough to watch.

03 Inside the Live Build

Four agents, built in front of the room in about thirty minutes, against a real Part 135 operation's live systems and live data. Each one was given a single job. All of them could reach each other in a shared workspace, which is what turned four assistants into something that behaved like a small team.



The inbox assistant

Identity · Own mailbox · Own credentials

Given agency to answer, and told to treat all mail and attachments as untrusted.

ANSWERED LIVE



The flight dispatcher

Authoritative sources only

Found the official weather APIs and FAA documentation, then read the METAR and TAF and converted UTC to local.

WROTE ITS OWN REFERENCE



The flight scheduler

Connected to the flight management system

Built its own query against the API, returned both aircraft and the open windows on request.

LIVE FLEET DATA



The charter sales assistant

Priced on rules, not on instinct

Checked availability with the scheduler, handed replies to the inbox assistant, and declined the trips that did not fit.

TALKED TO THE OTHERS

None of the four were clever. Each of them was narrow, sourced and scoped.

What Made Them Work

The interesting part of the demonstration was not the model. It was six choices, any of which an operation can copy this month.



One job each

An agent idle most of the month is a focused task handler, not waste.



Scoped credentials

Application passwords named to the bot, revocable in one click.



Untrusted by default

An instruction arriving by email is not a command. Injection attempts failed.



Its own machine

When no API exists it opens a browser, and can be taught by demonstration.



Cheap at rest

A scheduler wakes it. No new work means no model consumption.



Email as the API

Where no integration exists, unstructured mail becomes a structured record.

HOW ONE OF THEM GETS BUILT

1

Name the single job

One function, described the way a new hire would be briefed.

2

Give it rules and sources

The pricing logic, the approved data, the documentation to cite.

3

Issue a scoped credential

Specific to that agent, limited in reach, destroyed the day trust ends.

4

Prove it in a sandbox

Run it beside the operation before it runs inside the operation.

What it costs to run

Token consumption, asked and answered in the room

- A scheduler, not the model, does the waiting and the polling
- Classification can read only the subject line, and the body only when in doubt
- Sending and receiving mail runs on the mail protocols, not on the model
- Consumption gets large when building software, not when running a business function

THE TAKEAWAY FROM THIRTY MINUTES

A capable agent is now a short exercise. That is the good news and the warning in the same sentence, because the constraint has moved from what can be built to what should be allowed to run.

The Data Layer Underneath

Every agent above is only as good as what it can reach. The frontier models already read the public internet. The advantage left to an operator is private, organized and accessible data, which is a content problem long before it is a technology one.



Markdown is the accessible standard

Write for machines, export for people.

Why: plain text with light structure is the cheapest thing for a model to read accurately, and more of it has been written in the last four years than in the rest of recorded history. **Do this:** rewrite the operations manual in Markdown and export the PDF as the official artifact. Scanned and photocopied documents remain the hardest case.



The private library is the moat

Retrieval over an owned corpus.

The example: a Georgia poultry operation spent two years filming equipment and interviewing maintenance workers of 50 years who insisted their work could not be written down. **The result:** proprietary knowledge no public model has, answering questions no public model can. That corpus, not the model, is the competitive asset.



Determinism where it counts

Imaginative by design is not always wanted.

The trap: these models are non-deterministic, which is a feature and the source of hallucination. Finance and operations need the same answer every time. **The pattern:** let AI interpret and standardize at the edges, where vocabulary varies, and let it build the deterministic system that runs in the middle.

"Garbage in, garbage out was always true. Agents make it expensive."

04 Authority, Not Enthusiasm

Capability is well ahead of governance, and the gap between them is where an operator gets hurt. The instrument offered in the session was a six level participation taxonomy: permission climbs the way capability climbs, and it climbs in writing.

1

Prohibited

The function is off limits to AI entirely, and that is stated.

2

Reference

It may look things up. A person does everything else.

3

Flag

It may surface what deserves attention, and stops there.

4

Draft

It may prepare work a named person reviews and sends.

5

Recommend

It may propose the decision. A person still makes it.

6

Automate under supervision

It acts, inside a defined scope, with oversight that is real.

THE RULE THAT MAKES IT USABLE

Classify the function, never the product. The inbox assistant reads a mailbox, decides what matters, replies to a customer, and reaches the scheduling system. Those are four functions, and they can sit at four different levels. Assigning one level to a tool, or to a vendor, is how the discipline quietly fails.

THE LESSON FROM THE DEMONSTRATION

A prompt is a request, not a boundary.

During the live build, one agent found an earlier connector on its own, found a client and a token, and connected without being asked. The instruction not to do that would have been a sentence someone remembered to type. **A boundary lives in a scoped credential and in written authority**, and nowhere else.

An unwritten level is still a level. It is just **being chosen by somebody else**.

The Operator's Acceptance Test

Four questions, asked of every function before it goes anywhere near an operation. The first is a matter of fact. The other three are decisions, and they belong to the operator.

Name the seat, not the tool.

Accountability

When the agent replies to a customer, who owns that reply

Every message that leaves the operation has an owner today. An agent does not change that requirement, it only makes the answer harder to locate after the fact.

In the ops manual, not in a chat.

Authorization

Who approved the connection, and where is that written

Someone approved every credential an agent holds. If that approval lives only in a password manager or a memory, the authority is informal, and informal authority does not survive an audit.

Reconstructable afterwards.

Evidence

Could it be shown what it read and why it answered that way

Weather is the clean example. If the operation specifications name the authorized source, the agent has to pull from that source, and the record has to show that it did.

Failure is usually quiet.

Silence

When it stops working and says nothing, who notices

Injected commands get the attention, but the quieter failure is the agent that simply stops. The question is not whether that happens. It is how many hours pass first.

WHAT "A HUMAN REVIEWS IT" ACTUALLY REQUIRES

Five parts, all of them: see it, understand it, be able to decide against it, have the time to do so, and be able to intervene. A reviewer missing any one of them is a passenger, not a crew member.

05 What to Do Next

None of this requires a platform decision this quarter. Three moves, in this order, put an operation meaningfully ahead of where the room started.

1

Write the levels down

About half the room had an AI policy. Fewer had one that named a level per function, or that ruled out personal and free accounts for company work.

2

Take the day back

Classification, routing rules and a morning brief are Gear 1 work, available now, and they are most of that day a week.

3

Build one narrow agent

One job, one scoped credential, one named owner, and a sandbox before it touches a live operation.

HIGH IMPACT · LOW EFFORT

Do first

Mail triage and routing, and converting the manuals to Markdown.

HIGH IMPACT · HIGH EFFORT

Plan for

The private library, and agents connected to operating systems.

LOW IMPACT · LOW EFFORT

Fill-ins

Meeting summaries, internal reporting, recurring status assembly.

LOW IMPACT · HIGH EFFORT

Avoid

One all purpose assistant holding broad credentials and no scope.

What a committed quarter looks like

The shape of Gear 2 inside an operation

- A leadership team that has each built and shared one working routine
- An AI policy that names a level for every function, not a verdict on a vendor
- One agent in live use, owned by a named seat, with a revocable credential
- The core manuals converted and accessible to the systems that need them

THE HONEST CAUTION

Security, privacy and quality control across a growing population of agents is the challenge arriving next, and it is worth watching now rather than buying into today. It is not a reason to wait, it is a reason to write the rules while the count is still small.

Agentic sprawl is next year's problem. Not starting is this year's.

Thank you for a genuinely curious room.

Every speaker is reachable, and all four are glad to walk any part of this through with a team.



Jason Rhoades

MEANINGFUL AI

jason@meaningful.ai



Diego Magrini

NERD.AERO

diego@nerd.aero



Jiri Marousek

AIRTERA

jiri.marousek@airtera.com



Toby Benenson, AIGP

FOUNDER · SAYFLIGHT

toby@sayflight.ai