

Beyond the Hype: Making AI Work in Business Aviation



Srikanth Sopirala

Principal GenAI/ML Specialist Solutions Architect, AWS

Suresh Raavi

Senior Solutions Architect, AWS

Quick Audience Check

BEFORE WE BEGIN

Before we dive in, a quick show of hands:

- 1 How many of you have used an AI tool — any AI tool — in the last 30 days?
- 2 How many have tried AI on a work task — not just asking ChatGPT a question, but actually applying it to a workflow?
- 3 **How many are actively piloting or deploying AI in your aviation operations today?**

Why this matters

Your answers shape how we spend the next 45 minutes. We'll meet you where you are.

Goal

By the end, every hand goes up for #2 — you'll leave with a concrete starting point.

Agenda

AI Foundations → 3 Use Cases → Adoption Playbook → Responsible AI → What's Next

Why AI, Why Now: The Pressure on Business Aviation

THE CHALLENGES DRIVING URGENCY

It's 2 AM. A Gulfstream G650 is grounded in Teterboro.

Your technician is searching 4,000 pages of manuals for one torque spec. Every minute costs \$2,500.

Meanwhile, your best client texts: "Where's my plane?" — and no one is awake to answer.

This is the reality AI solves today. Here are the four pressures driving urgency:



\$10K–150K

per hour

Grounded Aircraft (AOG)

Every hour an aircraft is grounded for maintenance costs \$10K–150K. Industry-wide, groundings and MRO inefficiency destroy ~\$25–27B in value every year.



35–40%

avoidable

Reactive Maintenance

Predictive maintenance can cut unscheduled events 35–40% and lift dispatch reliability from 97.5% to 99.2% — yet most answers still sit in thousands of manual pages.



728,000

techs short

Workforce Cliff

The industry needs 728,000 new technicians in 20 years — with a 20% shortfall projected by 2028 and ~30% of expertise lost each time a veteran retires.



~11,500

flights/day

Rising Expectations

With ~11,500 private-jet flights a day, clients now expect instant, 24/7, concierge-level responses — a service load humans alone can't sustain.



01

AI Foundations for Aviation

Demystifying the AI stack — and leveling-set on the terminology we'll use all session.

The AI Stack in 60 Seconds

FOUR BUILDING BLOCKS — THE ONLY ONES YOU NEED TODAY



Predictive ML

Learns from historical data to forecast. →
"Will this part fail in the next 100 hours?"




Generative AI + RAG

Reads, summarizes, answers from your
documents — with citations.



Foundation Models (Amazon Bedrock)

Pre-trained intelligence as a service — no
ML team required.



AI Agents

Models that take action — call APIs,
orchestrate multi-step workflows.

Where AI Delivers Value in Aviation Today

THE OPPORTUNITY LANDSCAPE



Operations

Maintenance knowledge, scheduling, predictive upkeep, and turnaround efficiency.

Use Case 1



Customer Engagement

Booking, concierge service, and 24/7 responsiveness across channels.

Use Case 2



Compliance

Documentation, filings, audit readiness, and consistency at scale.

Use Case 3



Safety & Knowledge

Capturing expert knowledge and surfacing it safely before it retires.

Cross-cutting

Why AWS for Aviation

THE PLATFORM BEHIND THE USE CASES

Why choose AWS? Four reasons that matter for aviation — plus published reference architectures so you're not starting from zero:



Choose

your model

Model Choice, Not Lock-In

Amazon Bedrock gives you access to Claude, Llama, Nova, Titan, and more. Pick the best model for each job — no single-vendor dependency.



Your

data stays yours

Data Sovereignty

Your data is never used to train models. Encrypted at rest and in transit. Runs in your VPC with full compliance controls.



Built-in

safety

Bedrock Guardrails

Filter hallucinations, enforce policy, and keep AI on-domain — critical for safety-regulated industries like aviation.



Agents

at scale

Amazon Bedrock AgentCore

Deploy, manage, and monitor AI agents in production with built-in identity, auth, observability, and memory — from prototype to fleet-wide operations.



02

Use Case 1: Maintenance Knowledge Assistant

Instant, cited answers from thousands of pages of manuals — powered by RAG.

Maintenance Knowledge: The Problem & Business Case

USE CASE 1 · WHY IT MATTERS

When an aircraft is grounded, every minute counts. Technicians search thousands of pages under intense AOG pressure — and the experts who know the answers are retiring.



\$10K–150K

per hour

Cost of AOG Downtime

Every hour grounded costs \$10K–150K. Industry-wide, groundings and MRO inefficiency destroy ~\$25–27B in value every year.



20 min

→ 1–2 min

Technician Productivity

Textron's TAMI assistant cut troubleshooting time across 60,000+ pages of docs — a pattern AWS ships as a reference architecture.



2nd

highest cost

Risk of Inconsistency

Maintenance is 10–15% of operating cost, after fuel. Inconsistent answers from manual searches add safety and compliance risk.



728,000

techs needed


The Workforce Cliff

728K new technicians needed over 20 years — ~2/3 replacing retirees. ~30% of operational knowledge lost each time an expert departs.


Architecture Deep Dive: RAG & Data Flow

RETRIEVAL-AUGMENTED GENERATION, STEP BY STEP

BUILD TIME — index your documents




QUERY TIME — answer a technician's question



Why it matters: RAG grounds the model in YOUR documents — so answers are accurate, current, and cite their source. Guardrails filter anything off-policy.

AWS Architecture: Maintenance Knowledge Assistant

REFERENCE ARCHITECTURE · RETRIEVAL-AUGMENTED GENERATION





03

Use Case 2: Intelligent Booking & Scheduling

Voice-driven charter booking, available 24/7 — powered by AI agents.

The Business Problem: Demand Never Sleeps

USE CASE 2 · BOOKING & SCHEDULING

Charter demand never sleeps — but staffing does.

Charter booking is genuinely complex: aircraft availability, crew duty limits, positioning flights, and dynamic pricing all interact.

Requests arrive 24/7 across time zones, but expert schedulers aren't always available to respond instantly.

Customer expectations are rising — clients want concierge-level, instant answers, not a callback tomorrow.

~11,500

private jet flights per day worldwide — demand that never sleeps

\$16–26B

private jet charter market — clients expect concierge-level, instant service

Why Voice, Why Now

USE CASE 2 · WHAT CHANGED



Voice used to be painful

- Rigid phone trees and keywords
- Robotic, high-latency responses
- Brittle — broke on anything unexpected
- Customers preferred to wait for a human




What changed

- Speech-to-speech models (Amazon Nova Sonic)
- Natural, low-latency, interruptible conversation
- Understands intent, not just keywords
- Always-on = capacity without adding headcount

Architecture Deep Dive: Agent Orchestration


HOW AN AI AGENT PLANS AND ACTS



Key idea: an agent doesn't just talk — it plans, calls tools, and acts within guardrails, escalating to humans when the stakes are high.

AWS Architecture: Intelligent Booking & Scheduling

REFERENCE ARCHITECTURE · CONVERSATIONAL VOICE AGENT



Outcome: What to Notice

USE CASE 2 · THE PAYOFF



Always-On Capacity

+10%

VistaJet saw a 10% rise in on-time departures/arrivals after adding AI to flight planning & dispatch.



Agent Productivity

+14–34%

Rigorous NBER research: 14% avg productivity gain per contact-center agent, up to 34% for newer staff.



Revenue Upside

up to +30%

Travel agencies using AI voice tools report up to 30% revenue lift via personalization & upsell.



04

Use Case 3: Compliance & Document Automation

Faster, more consistent FAA documentation — with a human in the loop.

The Business Problem: Manual, Repetitive, Risky

USE CASE 3 • COMPLIANCE

Compliance work is high-stakes — and painfully manual.

FAA compliance documentation is manual, repetitive, and error-prone.

Staff re-key the same data across forms, filings, and logs — and small errors carry real regulatory risk.

Audits demand consistency and a clear paper trail that's hard to maintain by hand.




The stakes

Regulatory findings, rework, and delayed operations — all flow from inconsistent or late documentation.

Architecture Deep Dive: Intelligent Document Processing


HOW AI READS, DRAFTS, AND ROUTES FOR REVIEW



Key idea: AI drafts, people decide — and the audit trail proves it. Consistency at scale, without giving up control.

AWS Architecture: Compliance & Document Automation

REFERENCE ARCHITECTURE · INTELLIGENT DOCUMENT PROCESSING



Outcome & Deeper Look: Accuracy + Auditability

USE CASE 3 · WHY YOU CAN TRUST IT



Consistency

The same rules applied to every filing — no drift between people or shifts.



Confidence Scoring

Low-confidence extractions are flagged for human attention automatically.



Validation Checks

Business rules verify completeness before anything reaches a reviewer.



Audit Trail

Every step is logged — source, draft, edit, approval — for clean audits.

Real proof: An FAA aircraft-records platform built on AWS (Claude on Bedrock + Textract + OpenSearch) cut a ~30-day conformity audit to about a third of that time at up to 96.6% classification accuracy — scaling toward 100K–200K pages/month with full citation traceability.



05

AI Adoption Playbook

Three examples — maintenance, booking, compliance. So how do YOU get started?

A Practical 3-Step Framework

START SMALL, PROVE VALUE, THEN SCALE

1

Find

Pick one painful, high-frequency task with a clear owner and measurable value. Resist the urge to transform everything at once.

2

Prove

Build a focused pilot on real data. Measure time saved, cost avoided, and quality. Let the numbers make the case.


3

Scale

Templatize what works, add guardrails, and expand to adjacent use cases with the lessons you learned.

The Quick-Win Matrix

VALUE VS. EFFORT — WHERE TO START



Common Pitfalls to Avoid

WHAT DERAILS AI ADOPTION



Boiling the Ocean

Trying to transform everything at once. Start with one use case and win.




No Success Metric

If you can't measure it, you can't scale it or defend the investment.



Ignoring Data Readiness

Garbage in, garbage out. Clean, accessible documents make or break RAG.



Skipping the Frontline

Adoption fails without the people who'll actually use it every day.

Your First 90 Days

THE MONDAY-MORNING ACTION PLAN

DAYS 0–30

Pick & Scope

- Choose one high-value use case
- Name an owner and success metric
- Confirm the data you'll need

DAYS 30–60

Build the Pilot

- Stand up a focused proof of concept
- Use real documents and real users
- Keep scope tight and time-boxed

DAYS 60–90

Measure & Decide

- Compare against your baseline
- Capture lessons and guardrails
- Decide: scale, iterate, or stop

The BREAK Mindset: Reimagine, Don't Just Automate

ASK 'HOW DO WE REIMAGINE THIS?' — NOT 'HOW DO WE AUTOMATE EXISTING STEPS?'

B

Blindspot Scan

What are we taking for granted?

✂ *Are paper-based processes still best, or just habit?*

R

Radicalize Constraints

Force radical redesign by imposing extreme constraints.

✂ *What if you ran ops with 50% fewer manual steps?*

E

Economic Dissection

Where are hidden costs and value-destroying elements?

✂ *What does a single AOG really cost across the entire chain?*

A

Assumption Audit

Which 'requirements' are actually just habits?

✂ *Does every booking need a human, or just the complex ones?*

K

Kaizen the Happy Path

Relentlessly optimize the end-to-end ideal workflow.

✂ *What does perfect maintenance-to-clearance look like?*

New ways to create value — not just faster execution of old ways. Source: AWS Enterprise Strategy.



Responsible AI & Trust

Adopting AI with confidence — trust, guardrails, and human oversight.

The Trust Gradient: Automate → Augment → Alert

MATCH AUTONOMY TO RISK

Automate

Low-risk, repetitive tasks. Let AI run with light oversight.

e.g. sorting & tagging incoming documents

Augment

AI drafts, people decide. The default for most use cases.

e.g. drafting a compliance filing for review

Alert

High-stakes calls. AI flags; humans own the decision.

e.g. anything affecting safety or airworthiness

Guardrails & Data Privacy

TRUST IS ENGINEERED IN



Bedrock Guardrails

Filter unsafe, off-topic, or non-compliant output before it ever reaches a user.



Source Citations

Answers point back to the source document — no unverifiable claims.



Data Stays Private

Your data remains in your AWS environment, encrypted and access-controlled.



Not Used for Training

Your content is not used to train the foundation models. Full stop.

Human-in-the-Loop: People Stay Accountable


AI ACCELERATES — IT DOESN'T DECIDE

The people stay accountable — AI just makes them faster.

People stay in the loop wherever safety, compliance, or high value is at stake.

AI accelerates the work — drafting, retrieving, summarizing — but it does not replace the human sign-off.

This is what builds durable trust with regulators, crews, and customers alike.





What's Next (12–18 Months)

Where business aviation AI is heading — and what to watch for.

The Near-Term Horizon

THREE SHIFTS TO WATCH



Multi-Agent Systems

Specialized agents coordinating end-to-end workflows — booking, dispatch, and maintenance agents talking to each other.

One request, many agents, one seamless result.



Multimodal AI

Voice, image, and document understood together — snap a photo of a part and get the procedure and the order.

See it, say it, solve it.



Aviation-Specific Models

Foundation models fine-tuned on aviation language, regulations, and operations for sharper, safer results.

Models that already speak aviation.

Key Takeaways

WHAT TO REMEMBER

1

AI delivers value in aviation today

Maintenance, booking, and compliance are proven, ready starting points — not someday.

2

Start with one focused, high-value use case

Prove it, measure it, then scale what works to adjacent problems.

3

Adopt responsibly

Guardrails and human-in-the-loop keep AI safe, trusted, and compliant.

AISBA 2026 · Beyond the Hype: Making AI Work in Business Aviation

The Technology Is Ready.

Are You?

Your competitors are already piloting. Will you lead — or react? → Let's scope your first pilot — come find us today.

Srikanth Sopirala

Principal GenAI/ML Specialist Solutions Architect, AWS

sopiral@amazon.com

Suresh Raavi

Senior Solutions Architect, AWS

sraavi@amazon.com

Sources & References (1 of 2)

MARKET, INDUSTRY & WORKFORCE DATA

1. **AI in Aviation Market (2035 outlook)** — GMIInsights gminsights.com/industry-analysis/artificial-intelligence-in-aviation-market
 2. **AI in Aviation Market Report 2025–30** — MarketsandMarkets marketsandmarkets.com/Market-Reports/artificial-intelligence-in-aviation-market-106037016.html
 3. **AI in Aviation Market to 2035** — Transparency Market Research transparencymarketresearch.com/ai-in-aviation-market.html
 4. **AI in Aviation Market 2026–33** — Grand View Research grandviewresearch.com/industry-analysis/ai-aviation-market-report
 5. **AI in Aviation Market Size & Share** — Fortune Business Insights fortunebusinessinsights.com/ai-in-aviation-market-113289
 6. **AI-Enabled Predictive Maintenance in Aerospace** — Fortune Business Insights fortunebusinessinsights.com/ai-enabled-predictive-maintenance-in-aerospace-market-117506
 7. **Aircraft Predictive Maintenance Market** — GMIInsights gminsights.com/industry-analysis/aircraft-predictive-maintenance-market
1. **Air Charter Services Market** — Fortune Business Insights fortunebusinessinsights.com/air-charter-service-market-108388
 2. **Private Jet Charter Services Market** — Mordor Intelligence mordorintelligence.com/industry-reports/private-jet-charter-services-market
 3. **Business Aviation Services Market** — Fortune Business Insights fortunebusinessinsights.com/business-aviation-services-market-118064
 4. **Global Business Jet Fleet (24,442)** — Airbus Corporate Jets acj.airbus.com/en/newsroom/stories/2025-06-study-...business-jet-fleet
 5. **2025 Business Aviation Market (3.89M departures)** — JETNET jetnet.com/resources/blog/the-2025-business-aviation-market...
 6. **AOG Cost per Hour (\$10K–150K, Boeing benchmark)** — Actionstreamer / AAA Air Support actionstreamer.com/blog/what-does-aog-mean-and-what-does-it-cost-airlines-per-hour
 7. **Older Aircraft, Too Few Technicians (\$25–27B/yr)** — BCG bcg.com/publications/2026/older-aircraft-too-few-technicians-the-aerospace-fix
 8. **Pilot & Technician Outlook (728K techs)** — Boeing boeing.com/commercial/market/pilot-technician-outlook

Sources & References (2 of 2)

CASE STUDIES & AWS REFERENCE ARCHITECTURES

1. **Textron Aviation TAMI (20min→1–2min, 60K pages)** — Microsoft Customer Story microsoft.com/en/customers/story/23024-textron-aviation-azure-open-ai-service
2. **VistaJet +10% on-time; PdM -20% events** — CARI Journals (JBSM) carijournals.org/journals/index.php/JBSM/article/download/2284/2691/6786
3. **Predictive Maintenance 35–40% fewer events** — NBAA nbaa.org/aircraft-operations/safety/.../how-trend-analysis-informs-predictive-aircraft-maintenance
4. **Flexjet + GE Digital Maintenance Insight** — Flexjet flexjet.com/press/flexjet-to-adopt-ge-digital-maintenance-insight-fleetwide
5. **FAA Compliance Automation (96.6%, ~30d→½)** — New Math Data (AWS case study) newmathdata.com/case-studies/aviation-ai-compliance-platform
6. **AI Agent Productivity +14–34%** — Brynjolfsson/Li/Raymond, NBER 2023 voiceaiwrapper.com/insights/measuring-roi-the-business-case-for-ai-in-your-call-center
7. **WilcoIQ FBO Voice AI** — WilcoIQ wilcoiq.com
8. **Infosys Aviation Maintenance on Bedrock** — Infosys infosys.com/industries/aerospace-defense/insights/transforming-aviation-maintenance.html

1. **Guidance: Aerospace Technician’s Assistant on AWS** — AWS Solutions aws.amazon.com/solutions/guidance/aerospace-technicians-assistant-on-aws
2. **Guidance: Aircraft Predictive Maintenance on AWS** — AWS Solutions aws.amazon.com/solutions/guidance/aircraft-predictive-maintenance-on-aws
3. **How Amazon Bedrock Knowledge Bases Work (RAG)** — AWS Bedrock User Guide docs.aws.amazon.com/bedrock/latest/userguide/kb-how-it-works.html
4. **Knowledge Bases for Amazon Bedrock (RAG options)** — AWS Prescriptive Guidance docs.aws.amazon.com/prescriptive-guidance/latest/retrieval-augmented-generation-options/rag-fully-managed-bedrock.html
5. **Intelligent Document Processing on AWS** — AWS Solutions docs.aws.amazon.com/solutions/intelligent-document-processing-on-aws
6. **Deploy a Full-Stack Voice AI Agent with Nova Sonic** — AWS ML Blog aws.amazon.com/blogs/machine-learning/deploy-a-full-stack-voice-ai-agent-with-amazon-nova-sonic
7. **How Amazon Bedrock Agents Works** — AWS Bedrock User Guide docs.aws.amazon.com/bedrock/latest/userguide/agents-how.html
8. **Aurora PostgreSQL (pgvector) as a Knowledge Base** — AWS Aurora User Guide docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.VectorDB.html

Note: Textron TAMI runs on Azure OpenAI (cited as an OEM benchmark). All architecture patterns above are AWS-native. Full 60+ source list retained in research notes.