WHITE PAPER

# The State of AI in Aviation

Assurance Before Autonomy: Navigating the Regulatory and Safety Gap

**Bilal A. Khan** | Fractional CTO — Aviation, Safety-Critical & AI Systems | July 2026


### Executive Summary

Aviation is entering its AI decade. On July 1, 2026, Jeppesen ForeFlight unveiled ForeFlight Airflow, an agentic AI engine its CEO, Brad Surak, framed not as artificial intelligence but as “aviation intelligence” — AI whose outputs are cross-checked against industry safety and governance protocols before they reach a human decision-maker. That framing captures the defining tension of this moment: the technology is ready faster than the certification framework built to trust it. This paper looks at where AI adoption in aviation stands today, why the industry’s foundational safety standards — DO-178B/C for software and DO-200A/B for aeronautical data — were never built for systems that learn, and what concrete options leaders have to move forward responsibly while regulators finish writing the rules.


## 1. Where AI Adoption Stands Today

The market signal is unambiguous. AI in aviation is projected to grow from roughly $1.76B in 2025 to $4.86B by 2030, a 22.6% CAGR, driven by predictive maintenance, flight planning, crew and fleet optimization, and air traffic trajectory prediction. ForeFlight Airflow is the clearest expression of where this is headed: an open, model-agnostic engine that lets operators plug in their own agents or third-party models while keeping “as much or as little human control over key decisions as their operations require.” Commercial and business aviation offerings are expected later in 2026, with military-focused capabilities to follow.

The pattern across the industry is consistent: AI is landing first in advisory and back-office roles — flight plan validation, maintenance prediction, dispatch, safety-trend analysis — and only gradually approaching flight-critical, deterministic control. That sequencing is not a marketing choice. It is a direct response to the certification gap described below.

## 2. The Core Tension: Deterministic Certification Meets Learning Systems

Two standards families are frequently conflated, but they solve different problems and share the same blind spot on AI:

- DO-178B/C — “Software Considerations in Airborne Systems and Equipment Certification,” the standard FAA, EASA, and Transport Canada use to approve virtually all commercial aerospace software. It assumes fully specified requirements, deterministic behavior, and requirements-based test coverage.
- DO-200A/B — “Standards for Processing Aeronautical Data,” governing the accuracy, traceability, timeliness, and integrity of the data pipeline (navigation databases, charts, procedures) from supplier to cockpit.

Machine learning breaks the core assumption both standards were built on: the same inputs can legitimately produce different outputs as a model learns, and behavior is defined statistically across a training distribution rather than exhaustively specified in advance. Neither DO-178C’s test-to-requirements model nor DO-200B’s data-lineage model was designed to assure a system whose “requirement” is a probability distribution learned from data. This is the gap the entire industry is now working to close — not by discarding these standards, but by building a companion layer of assurance around them.

| Standard / Body | Scope | AI-Era Status |
|---|---|---|
| DO-178B/C (RTCA/EUROCAE) | Airborne software design assurance (DAL A–E) | Deterministic model doesn't map to learned behavior |
| DO-200A/B (RTCA) | Aeronautical data quality, traceability, integrity | Principles now being extended to training/validation data |
| FAA Roadmap for AI Safety Assurance | Agency-wide guiding principles for AI in aircraft & operations | Living document; certification position papers targeted Q1 2026 |
| EASA AI Roadmap 2.0 / Concept Paper | Human-centric AI guidance by autonomy level (1–3) | Level 1–2 guidance published; Level 3 in progress; policy by 2028 |
| EUROCAE WG-114 / SAE G-34 (AS6983 / ED-324) | Certification process standard for AI-based aeronautical products | 600+ participants incl. Boeing, Airbus, FAA, EASA; in development |

## 3. The Regulatory Response: Incremental, Not Reactive

Regulators are not stalling; they are sequencing. The FAA's Roadmap for AI Safety Assurance sets out an incremental, safety-first path: deploy AI in low-risk environments first, gather evidence, and expand scope only as confidence is established — while pursuing global harmonization with EASA and other authorities. Its Overarching Properties initiative, running since 2016, offers a compliance path that does not require rewriting DO-178C: a system can be certified by demonstrating Intent (the function is correct and complete), Correctness (the implementation meets that intent under foreseeable conditions), and Innocuity (anything beyond the intended function causes no unacceptable safety impact) — properties that can, in principle, be demonstrated for ML components without full behavioral determinism.

EASA has taken a parallel, level-based approach. Its AI Roadmap 2.0 and Concept Paper Issue 2 provide guidance for Level 1 (human augmentation) and Level 2 (human-AI cooperation) applications now, with Level 3 (advanced automation) guidance following and a consolidated AI/ML policy targeted for 2028. Underpinning both agencies is the joint EUROCAE WG-114 / SAE G-34 working group — more than 600 participants from Boeing, Airbus, the FAA, EASA, and others — developing AS6983/ED-324, a dedicated certification process standard for AI-based aeronautical products, with its first scope deliberately limited to frozen models trained via supervised learning. The message from every one of these efforts is the same: assurance before autonomy, and narrow scope before broad scope.

## 4. Options to Close the Gap

Leaders don't need to wait for finished standards to move responsibly. Five approaches are converging across the industry:

- Sequence by risk, not by ambition. Deploy AI first in advisory, ground-based, and predictive-maintenance roles — exactly where ForeFlight Airflow and the FAA's incremental approach both start — before touching flight-critical control.
- Apply run-time assurance (RTA). Rather than proving an ML component is deterministic, wrap it with a certifiable monitor that bounds its outputs and reverts to a safe, verified fallback. This is broadly seen as the most viable near-term path to fielding advanced ML in safety-critical loops.

- Extend DO-200B discipline to training data. The same rigor — accuracy, traceability, completeness, timeliness — that governs navigation data today is now the leading model for assuring the datasets that train and validate ML systems, regardless of which final standard lands.
- Design for explainability and human oversight deliberately. Regulators consistently pair AI assurance with human-factors requirements; model-agnostic, adjustable-autonomy architectures (again, Airflow is a live example) satisfy this by design rather than as an afterthought.
- Get a seat at the standards table. AS6983/ED-324 and the FAA/EASA policy frameworks are being written now, not someday. Organizations participating in WG-114/G-34 or responding to EASA/FAA consultations are shaping what “compliant” will mean for everyone else in 2027–2028.

## 5. What This Means for Engineering Leaders

None of this is unique to avionics primes. Any organization building AI into a regulated, safety-critical, or high-consequence workflow — aviation, industrial IoT, medical devices, critical infrastructure — is facing a version of the same problem: proven governance frameworks built for deterministic systems, meeting AI that behaves probabilistically. The organizations that will move fastest over the next three years are not the ones with the most advanced models. They are the ones that started building data governance, run-time monitoring, and human-oversight discipline before the standards forced them to.

---

*Where does your organization’s AI roadmap sit relative to where regulators are heading — ahead, aligned, or exposed?*

I help engineering and safety leaders in aviation and other safety-critical industries build pragmatic, compliance-ready AI adoption plans — grounded in the same standards and certification realities covered here. If you’re weighing how and where to start, I’d welcome the conversation.

**Bilal A. Khan | Fractional CTO | bilal.khan@stragentech.com**

## Sources


- Jeppesen ForeFlight — Agentic AI Strategy and Engine announcement (July 1, 2026)

globenewswire.com/news-release/2026/07/01/3320895

- ForeFlight Airflow — early access

ai.jeppesenforeflight.com

- FAA — Roadmap for Artificial Intelligence Safety Assurance

faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance

- FAA — Overarching Properties initiative (via INCOSE / Wiley)

incose.onlinelibrary.wiley.com/doi/10.1002/iis2.13111

- EASA — Artificial Intelligence Roadmap 2.0

easa.europa.eu — AI Roadmap 2.0

- EASA — Artificial Intelligence Concept Paper Issue 2 (Level 1 & 2 ML applications)

easa.europa.eu — AI Concept Paper Issue 2

- EUROCAE WG-114 / SAE G-34 — AS6983 process standard for AI certification

standict.eu — EUROCAE WG114/SAE G34 AS6983

- DO-178C overview

en.wikipedia.org/wiki/DO-178C

- DO-200B — Standards for Processing Aeronautical Data

FAA AC 20-153B

- AI in Aviation market sizing 2025–2030

MarketsandMarkets — AI in Aviation Market Report

- Run-time assurance for safety-critical control systems

arXiv:2110.03506