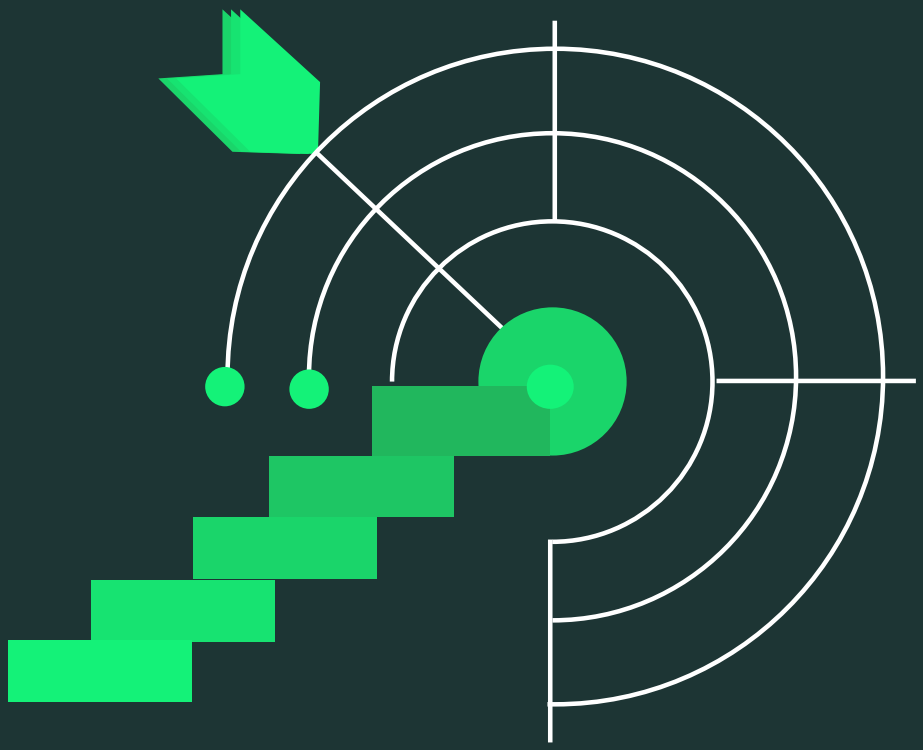

Runway to revenue

The strategic value of data and AI in aviation maintenance and operations

kubrick

Contents

kubrick

Executive summary

While demand surges, the aviation industry is being stretched beyond capacity, due to aircraft and technician shortages, rising operational costs, and increasing customer expectations. One of the strategic levers that turns revenue into profit is maintenance and operations.

AI shows the promise to address supply chain and talent bottlenecks cascading since the pandemic and help airlines recover lost revenue, turning data into the aviation sector's second most important asset after fleets
and airports.

This report explores key use cases for data and AI across the maintenance and operations lifecycle that minimizes cost and builds resilience, creating the runway to drive more revenue.

Research approach

This report combines senior stakeholder and market research with Kubrick's experience of enabling AI capability in aviation teams in world-leading airlines. The discussion is guided by insights provided by a selection of data leaders in the sector working in partnership with Databricks. The select body of leaders represent aviation organizations including JetBlue, United Airlines, Delta, and Heathrow.

Introduction

Maintenance and operations as a critical strategic lever

The profitability challenge

Post-pandemic demand is soaring. In 2024, air travel exceeded pre-pandemic levels for the first time to reach 5 billion passengers. 2025 topped it, growing by another 4.4% - and 2026 is projected to be an even more intensive year, carrying 5.2 billion passengers.

Despite skyrocketing demand, net profit margins only increased 0.3% from 2024 to 2025[1] due to factors such as unstable fuel costs, supply chain disruption, and weather, which make translating revenue into profitability an existential challenge.

With tightening margins, airlines aim to seize market share of passenger and cargo demand, taking on ambitious schedules and routes, while maintaining the quality of service. To do so, they need fleets and services that can operate at maximum capacity – or risk turning revenue into loss. A flight cancelled due to a maintenance or scheduling issue can cost anywhere from $40,000-$125,000, depending on its size, between lost seats, rebooked tickets, compensation payouts, crew and labor costs, and other knock-on effects (baggage, customer service).

The airlines that stretched themselves have faced greater consequences than delay and cancellation compensations. In 2025, the US Department of Transport issued record fines of over $4m for tarmac delays and begun issuing fines for 'chronic delays' of $2m[2].

In the hypercompetitive aviation market, the role of maintenance and operations is key profitability driver.

Taking Maintenance and Operations up a gear

An airline is only strong as its fleet. The aircraft manufacturing supply chain continues to face severe post-pandemic delays, leading to a global shortage of an estimated 5,000+[3], as well as many critical parts and materials for maintenance. Most notably, the Boeing-777Xs is now scheduled for delivery in 2027[4], seven years later than planned.

The focus is on making aging fleets last longer to keep up with their demanding schedules. In fact, the average age of a plane in service is now 15 years, compared to 13 years pre-pandemic[5]. Akin to the aging aircrafts, maintenance operations are growing outdated in the digital world. Technicians must search through thousands of pieces of written documentation, many of which are still on paper: OEM maintenance manuals (AMM), illustrated parts catalogs (IPC), fault isolation manuals (FIM), wiring diagrams (WDM), service bulletins (SBs), and historical logbooks and work cards.

While this information has been trapped in paper work cards and disjointed IT systems, modern data and AI techniques present a new way to interpret and use this data. The challenge for aviation is to turn data into their next best asset, transforming maintenance and operations into a strategic lever for resilience, expansion, and customer experience.

1

Maintenance

Creating AI-enabled technicians

The capability challenge: A technician shortage

The COVID-19 pandemic left the aviation industry with a maintenance talent gap that's still widening. When borders closed and fleets were grounded, thousands of experienced engineers, planners, and shop-floor technicians left the sector, and their deep industry knowledge went with them. According to one expert at StandardAero, organizations like Delta TechOps lost as much as 30,000 years of cumulative experience[6]. It is estimated that the industry will be short of 70,000 technicians by 2033[7].

Carriers are trying to rebuild their staff, with organizations like American Airlines opening 500 new technician roles in 2024 and 2025, and JetBlue announcing new partnerships with FAA-approved engineering colleges to train and onboard technicians directly into their workforce[8]. However, training junior staff to diagnose and fix faults confidently takes years and the gap between recruitment and retirement is widening.

The intermediate cracks are clear: delays and cancellations caused by issues within the airline's control have climbed from 5.2% in 2018 to 7.6% in 2023[9].

The new aircraft models are more complex, while aging fleets require consistent maintenance work to address chronic defects. Therein lies an opportunity for AI: to enable technicians to upskill quickly, keeping aging fleets running and anticipating the complexities of fleet modernization.

Troubleshooting: A data challenge

Most troubleshooting knowledge isn't neatly structured; it lives scattered across OEM manuals, handwritten work cards, old emails, chat logs, parts memos, and thousands of free-text defect reports, each with its own acronyms, shorthand, or typos.

To bridge this, many carriers have invested heavily in maintenance databases, which contain vast collections of internal troubleshooting notes, work card tweaks, engineering bulletins, and lessons learned. For some carriers, these knowledge bases are decades in the making and represent a significant competitive advantage.

But there's a catch: these systems are hard to keep current. They often live in disconnected silos, and difficult to search at speed. Carriers report taking years to get their data into a unified format to be able to use the databases or SaaS products available on the market for maintenance analytics. For a new technician or planner joining, the learning curve is steep; it can take years to fully absorb what's buried in those files.

The real value of this data is intelligent prioritization: using AI to read through the data, interpret spelling quirks and acronyms, and return only what matters for this aircraft, this scenario, this shift.

Use case: AI-assisted defect resolution

Every maintenance decision happens in the context of a tight turnaround. There's no time to scroll through PDFs at 2 am, when a cargo door fault is putting a high-value shipment at risk. Surfacing raw information isn't enough; technicians need the answer, matched to specific tail and the flight schedule.

Modern data cloud platforms provide the capability to ingest, transform, and unify data sources from across the MRO lifecycle: OEM manuals, maintenance logs and troubleshooting notes, work orders, flight schedules, inventory, and live flight data.

The release of Large Language Models (LLM) and Generative AI enables them to interrogate that data, providing technicians the opportunity to query in natural language for how to diagnose and resolve a defect.

For example, a technician can enter a fault code with a tail number into an app. An LLM, fine-tuned to the nuances of the aircraft model, can retrieve the most likely diagnosis and resolution, combining manual information and historic logs to recommend what has been most successful for that tail number before. With flight schedules and available inventory, it can prioritize the fixes on urgency and part availability.

Watch the demo here.

Case study: Saving $14m for a major US Airline

A global airline consistently ranked below competitors for on-time performance and faced regulatory fines. They rooted the cause to ECAM-related delays, where the junior technicians working the night shift struggled to make the right fixes in the allotted time.

In partnership with Databricks, Kubrick created the Defect Resolution AI Assistant to give technicians instant recommendations. This approach cut troubleshooting time by up to 50%, helping junior technicians handle the same complex tasks as their senior peers. This 'AI-memory' of fixes formed the basis of chronic defect identification, which has helped achieve a 30% reduction of maintenance-related delays, saving $14million in recovered revenue and customer compensation payouts.

Use case: Critical tail swaps

Maintenance and operations crews play a critical role in preventing delays and cancellations from cascading; every plane that goes out of service risks disrupting more downstream. However, conducting tail swaps to remove planes with potentially unresolvable defects is a challenge most often left to last-minute instinctive decisions. In major carriers, operations teams can make 50,000 swaps a month[10]. The unforeseen impact of these swaps can cause more complicated, costly impacts with aircraft and crew requirements/availability.

With advanced analytics, airlines can create digital twin simulations to give their planners foresight and predictability. Combining data like schedules, MEL status, parts, and crew availability, planners get a view on areas of congestion and idleness. When a tail number goes out-of-service, they can quickly simulate the downstream effects on the schedule, calculating the most efficient swaps to optimize on delay times, cost, customer experience, and overall disruption to the network.

Watch the demo here.

Case study: Preventing cascading delays with critical tail swaps

In a leading cargo carrier, planners were making last-minute swaps which resulted in dozens of downstream disruptions. A single event caused 62 cascading swaps which amounted to a 4-hour delay across the network at a cost of $100,000.

Kubrick developed a Critical Tail Swaps Simulator to tackle the issue. Built with the Databricks Data Intelligence platform and interfaced with an interactive dashboard system, showed planners how to reduce that scenario to just three swaps which limited disruption and cost. In production, the tool is shown to reduce cost of delays from Out-of-Service events by 10–18%, equating to $15-40m in savings in mid to large size carriers.

2

Planning

Recovering cost at scale

Creating resilience to the unstable MRO and fuel costs

Rising MRO costs persist as a key concern for aviation executives, as material shortages continue to inflate prices[11], making it the third highest outlay on the balance sheet after fuel and labor costs. While real-time troubleshooting helps keep flight on schedule, maintenance planning is strategic enabler for cost reduction and reliability that improves margins at scale.

When maintenance operations are scattered across national and international hubs, maintenance planning relies on a complex web of inventory and crew availability. Thus, identifying chronic defects enables inventory optimization for the right parts, which is critical to help overcome supply constraints.

Tackling the fuel problem

As the highest and volatile expense, fuel usage can make the difference between profitability and loss. Aviation emissions are also increasingly scrutinized, with targets to bring down the fuel consumption and carbon footprint.

Certain fuel levels are measured by regulatory standards and legally required, allotting for take-off and landing, taxiing, and alternate and contingency fuels for any change of plan or route, in addition to the flight fuel. However, other aspects are often determined by human judgment, such as arrival delay fuel allowance.

In these instances, even the smallest amount of unnecessary fuel leads to wastage; aircraft must burn extra fuel just to support the weight of the fuel they are carrying.

Case study: Fuel optimization

A major UK carrier wanted to investigate the wastage on their arrival delay fuel. Extreme instances of mis-anticipating delays had resulted in notable losses. For example, when a flight that had been assumed to have a 90-minute delay arrived on-time, the extra 8,550kg of extra fuel resulted in wasting $2,384 of fuel.

Combining key datasets, from live flight data, weather, and historic patterns, Kubrick created an advanced data science model to predict delay times based on weather conditions, time of day, and aircraft model. Tested on the flight data of 2024, Kubrick's model was shown to unlock $600,000-$1,600m in savings in just one area of fuel optimization.

Case study: Maintenance planning

A global airline was spending over $70m a year simply moving parts around their international network to conduct maintenance in different hubs where inventory was limited. They used Kubrick's advanced analytics capability to create a digital twin simulator to help anticipate routine maintenance with parts availability as tail numbers moved across the network.

With just a 10% reduction in wasted parts transportation, the airline saved an additional $7m in addition to the nightly maintenance savings done with the Defect Resolution AI Assistant.

3

Revenue

Maximizing opportunity

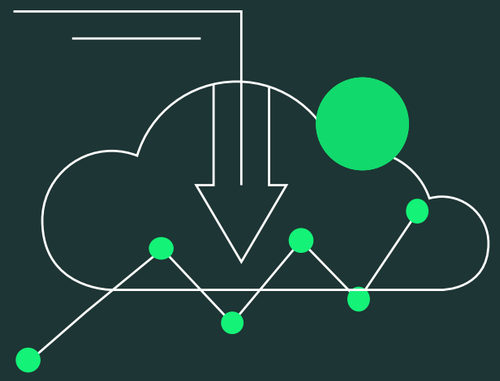

Maintenance pays off

On-time performance is not just about minimizing risk of fines; it's about competitive advantage. On-time performance correlates strongly with financial performance, with Return on Invested Capital (ROIC) doubling from 6% to 12% for airlines that arrive on time 70% of the time versus 90%[12].

On-time performance minimizes revenue loss of compensation payouts and missed seats on delayed and cancelled flights, but it is also a reflection of strong customer experience. Consumer surveys find that reliability and on-time performance is of increasing importance to customers, where airlines can no longer rely on loyalty schemes[13]. Now, maintenance is a key strategic lever for brand reputation and driving revenue through customer retention.

As fleet modernization becomes intrinsic to customer experience, so too will maintenance. Airlines now have a lot more to maintain than the aircraft: entertainment systems, reclining chairs, Wi-Fi service, electronically dimmable blinds, all require specialized maintenance.

These premium services are key competitive differentiators. However, they come with varying degrees of urgency when it comes to impact on revenue; some will be an inconvenience, some will 'lock out' a seat, some will delay an entire plane – or mark it out of service. Fleet modernization means technicians need more data to interpret to understand the defects, resolutions, and prioritization.

Use case: Tackling chronic defects to maximize revenue

Up against tight margins, airlines focus on first-class and business cabins. One major airline reported a loss of $1.2m in revenue from 'locked out' seats, where chronic defects reduced cabin capacity. However, identifying the critical fixes that keep seats in use amongst the noise of other minor defects, is a challenge.

Kubrick's aviation specialists developed an insights engine to detect the chronic defects which consistently lead to revenue loss. Using a clustering model, the engine clusters the myriad fault recordings on relevance and on frequency. In doing so, they were able to identify the 'occasional, 'repetitive', and 'chronic' defects that occur throughout flight cycles, providing insight to maintenance planners to address the most pressing issues before they even occur.

Concluding thoughts

The future of maintenance and operations for aviation

For decades, airline carriers have been tracking their operational footprint in disjointed, manual systems. Now, modern data platforms can unify these sources to create a blueprint for efficiency, cost reduction, and revenue uplift. Generative AI is making these insights actionable in real-time, right into the hands of frontline staff.

Airlines don't have to wait to strengthen their data foundations to start realizing value from AI. Tools like Kubrick's Defect Resolution AI Assistant show how to build these systems symbiotically, delivering value in a little as 90 days.

The runway to revenue is clear: the data is here, the technology is ready, and the opportunity is now. Those who act decisively will transform maintenance and operations from a cost center into a strategic engine for growth. The goal is something greater than the sum of its parts: every out-of-service event prevented by enhanced maintenance and operations equates to greater customer satisfaction and loyalty.

About Kubrick

Kubrick exists to transform lives through data & AI. We help global organizations realize lasting value from data and AI with a workforce we build ourselves, specializing in Aviation.

We deliver data and AI solutions that minimize operational cost, strengthen resilience against risk, and uncover revenue opportunity. Our clients can retain our people to drive lasting adoption while futureproofing their workforce with exceptional talent.

Since 2016, we've created over 3,000 data & AI specialists by removing the systemic barriers to the tech industry. We find incredible minds from all backgrounds to train with us, creating a diverse team of experts. We're the preferred partner of today's leading technology providers including Databricks to accelerate delivery and co-create revolutionary solutions.

Learn more about our data and AI solutions and specialists for airlines maintenance.

speaktous@kubrickgroup.com

References

[1] IATA

[2] US Department for Transport

[3] IATA

[4] Simply Flying

[5] Aerospace Global News

[6] Source: Expert Insight via AlphaSense

[7] McKinsey

[8] Business Wire

[9] CBS News

[10] Expert insight via Kubrick

[11] OliverWyman

[12] McKinsey

[13] OAG

kubrick www.kubrickgroup.com | email speaktous@kubrickgroup.com