

National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



Auditable AI Decision Intelligence for Aviation MRO: A KPI Governance Architecture

SeyyedAbdolHojjat MoghadasNian*

1. Tarbiat Modares University, s14110213@gmail.com


Abstract

Aviation Maintenance, Repair and Overhaul (MRO) organizations increasingly possess enterprise resource planning data, inventory records, work-order histories, procurement evidence, quality documentation, finance approvals, and customer commitments, yet many operational decisions remain delayed, fragmented, and weakly auditable. This article develops the Aviation Maintenance, Repair and Overhaul Decision Intelligence Governance Framework (AMRO-DIGF), a design-science architecture for governing artificial intelligence-enabled decision intelligence in safety-critical aviation maintenance environments. The study synthesizes aviation digital-operations guidance, airworthiness and repair-station obligations, AI risk-management standards, design-science theory, aircraft-maintenance AI literature, and a prior author-developed corpus on KPI-governed airline logistics, AISA-L, MRO efficiency, AI performance measurement, stock control, and predictive maintenance. AMRO-DIGF consists of five interdependent layers: data integration and lineage, operational diagnostics, AI-assisted recommendation, human authority and compliance governance, and KPI-based feedback. The contribution is not another dashboard or predictive-maintenance tool; it is a governance architecture that converts fragmented MRO evidence into traceable, human-authorized, compliance-aware, financially disciplined decisions. The article argues that AI creates credible MRO value only when recommendations are connected to evidence quality, bottleneck diagnosis, decision rights, airworthiness boundaries, override records, cybersecurity, and post-decision learning. The paper offers formula-level KPI logic for turnaround risk, bottleneck severity, parts readiness, margin leakage, AI recommendation quality, and governance compliance. It concludes that AMRO-DIGF should be evaluated through digital-twin simulation, expert review, and longitudinal case comparison before any claim of causal performance improvement is made.

Key words: Aviation MRO, Decision Intelligence, KPI Governance, Artificial Intelligence, Airworthiness Compliance, ERP, Maintenance Analytics

1. Introduction

Aviation maintenance, repair and overhaul is a high-consequence operating domain in which safety, fleet availability, spare-parts readiness, technician capacity, supplier reliability, customer commitments, and financial discipline must be coordinated under severe time



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



pressure. IATA positions digital aircraft operations across maintenance, supply chain, logistics, flight operations, ground operations, and aircraft asset transfer, indicating that maintenance decisions increasingly belong to a connected digital operating environment rather than to an isolated engineering function [1]. In this environment, enterprise resource planning systems, maintenance platforms, aircraft health monitoring, procurement workflows, inventory records, quality holds, finance approvals, and customer-scope changes generate large volumes of evidence but not automatically better decisions.

Artificial intelligence has entered this environment as a potential decision-intelligence layer. EASA frames aviation AI through a human-centric logic that connects safety, ethics, trustworthiness, assurance, and aviation certification [2]. FAA repair-station requirements and EASA maintenance-organization rules reinforce a similar boundary condition: maintenance decisions must remain grounded in controlled technical data, competent personnel, approved procedures, quality systems, traceable records, capability control, and authorized release-to-service accountability [3,4]. Therefore, AI in aviation maintenance cannot be treated as unconstrained automation. It must be governed as a decision-support capability inside an airworthiness accountability system.

The core problem is not the absence of digital systems. Aviation aftermarket platforms already support maintenance, parts logistics, manufacturing, aircraft services, and transactional control across work orders, inventory, purchasing, finance, and compliance evidence [7]. The unresolved challenge is that transaction data does not automatically become a defensible decision. A defensible maintenance decision must identify the binding operational issue, preserve evidence lineage, respect technical and regulatory constraints, assign decision ownership, record overrides, and learn from the outcome. NIST AI RMF and ISO/IEC 42001 strengthen this logic by treating trustworthy AI as a managed risk-and-accountability system rather than as isolated model performance [5,6].

This article addresses the gap between aviation maintenance data availability and governed decision execution. A delayed repair order or late aircraft release may arise from missing parts, tooling unavailability, engineering hold, technician-capacity limits, inspection delay, documentation gaps, vendor lead-time volatility, finance approval delay, or customer-scope ambiguity. Conventional dashboards show variance, but they often do not identify the binding bottleneck, rank feasible interventions, quantify margin leakage, or preserve a complete decision trail. The primary research question is: How can AI-enabled decision intelligence be governed and operationalized in aviation maintenance to reduce turnaround delays, diagnose bottlenecks, protect compliance, and improve financial discipline?

The article contributes to theory by repositioning aviation ERP and maintenance data infrastructures as decision-governance assets rather than only systems of record. It contributes methodologically by specifying a design-science artifact with layers, principles, KPI logic, and evaluation criteria. It contributes practically by giving airline technical operations, maintenance organizations, quality leaders, finance managers, and digital transformation teams a structured route from descriptive analytics to auditable decision execution. Its claim is deliberately conservative: AMRO-DIGF is a proposed governance architecture and validation pathway, not proof of causal improvement from a live proprietary deployment.




2. Literature Review

Design-science research is appropriate because the objective is to construct and justify an artifact for a real organizational problem. Hevner et al. define design science around the creation and evaluation of IT artifacts that respond to relevant business needs [8]. Peffers et al. translate this into a process of problem identification, objective definition, design and development, demonstration, evaluation, and communication [9]. Gregor and Hevner add that design-science contribution must clarify the artifact, its knowledge value, and its relation to practical relevance [10]. AMRO-DIGF is therefore assessed as a governance artifact: a structured architecture for converting maintenance evidence into accountable decisions.

Aircraft-maintenance AI literature establishes the operational relevance of prediction, diagnosis, optimization, and maintenance planning. Pavlyuk and Alomar review decades of AI applications in aircraft maintenance and show that the field covers aircraft health monitoring, airworthiness management, fault diagnosis, predictive maintenance, and maintenance operations [11]. Kabashkin et al. propose AI-driven fault detection and maintenance optimization for aviation technical support systems, while Shukla and Bansal connect AI prognostics to fleet management, operational efficiency, return on investment, downtime reduction, audit readiness, and compliance preparedness [12,13]. These studies show that AI is materially relevant to aviation maintenance, but they do not fully specify how AI outputs should be embedded in ERP-based decision rights, compliance escalation, quality gates, and financial-control logic.

AISA-L is the most direct antecedent of the present article. It develops a four-layer agentic AI strategy architecture that converts airline logistics KPIs into autonomous, auditable, sustainability-aligned decisions through perception, cognition, strategy, and action layers [14]. A related AI strategy-agent study also advances real-time KPI orchestration, ethical auditing, explainability, sustainability scoring, and disruption response in airline logistics [15]. These works move beyond static dashboards by treating KPIs as control variables within a closed governance-optimization cycle. The present article narrows and hardens that logic for aviation maintenance by adding airworthiness-constrained decision rights, repair-station accountability, quality-release boundaries, margin-leakage controls, and maintenance-specific bottleneck diagnosis.

Prior KPI-driven MRO research identifies aircraft availability, mean time between maintenance events, parts availability, safety, quality, customer satisfaction, and financial health as central to maintenance effectiveness [16]. A 360-degree airline logistics AI performance framework further shows that AI performance measurement must include safety, sustainability, efficiency, ethical AI, innovation, governance, and cross-functional alignment [17]. Stock-control and inventory-optimization studies show how AI can support spare-parts forecasting, replenishment automation, real-time analytics, and supplier coordination, all of which directly shape maintenance turnaround risk [18,19]. Predictive-maintenance and logistics digital-transformation research adds the need to integrate machine learning, IoT, ERP and maintenance systems, blockchain-enabled traceability, and data governance into operational reliability [20-22].



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



The gap is therefore precise. Existing literature explains AI prediction, KPI management, digital logistics, ERP modernization, and maintenance analytics, but it under-specifies maintenance decision execution as a governed chain from evidence to recommendation, human authority, compliance boundary, action, audit trail, and KPI feedback. AMRO-DIGF addresses this gap as a higher-order development of prior KPI and agentic AI work: it converts logistics-wide orchestration and maintenance KPI thinking into an aviation maintenance-specific decision-governance architecture.

3. Methodology

The study uses a conceptual design-science methodology. The problem-identification stage defines persistent decision failures in maintenance: fragmented evidence, delayed bottleneck diagnosis, uncertain decision ownership, weak auditability, and margin leakage. The objective-definition stage specifies the required artifact: a governance architecture that can convert operational evidence into AI-supported, human-authorized, compliance-aware decisions. The design-and-development stage constructs AMRO-DIGF through five layers, design principles, KPI constructs, and implementation controls. The demonstration stage is conceptual and scenario-based, focusing on turnaround delay, parts readiness, quality hold, supplier delay, and margin-risk decisions. The evaluation stage is proposed rather than claimed: expert review, digital-twin simulation, retrospective case comparison, and longitudinal implementation assessment.

The unit of analysis is the maintenance decision episode. A decision episode begins when an operational signal emerges, such as a late work order, missing part, quality hold, supplier delay, or margin-risk alert. It ends when the organization records a decision, assigns responsibility, authorizes action, preserves evidence, and evaluates the result. This unit is narrower than an airline, broader than a model output, and more useful than an isolated KPI value because it captures the full governance chain from evidence to action.

The article does not use confidential operational records, personal data, or proprietary airline datasets. It synthesizes public guidance, standards, peer-reviewed and preprint literature, and author-owned prior scholarly work as design inputs. This boundary protects against false empiricism. The paper does not claim statistical significance, implementation success, or generalizable causal performance improvement. Instead, it offers a validated-by-design artifact that can be tested in future empirical settings.

4. AMRO-DIGF Architecture

AMRO-DIGF consists of five interdependent layers. The first layer is data integration and lineage. Its purpose is to connect ERP, maintenance, inventory, procurement, quality, finance, customer, supplier, and aircraft-health evidence into a governed decision record. The essential controls are master-data quality, part-number normalization, work-order lineage, supplier-event timestamping, document completeness, access control, and change history. Without this layer, AI recommendations risk being technically persuasive but evidentially weak.



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



کد اختصاصی نمایه ISC
۰۴۲۵۰-۲۲۷۹۹



The second layer is operational diagnostics. It converts raw signals into a ranked explanation of the binding constraint. A late release may appear as a turnaround problem, but its true cause may be parts non-readiness, technician capacity, tooling, engineering disposition, quality inspection, supplier delay, customs clearance, or finance approval. The diagnostic layer should distinguish symptoms from constraints and separate controllable from non-controllable causes.

The third layer is AI-assisted recommendation. The recommendation layer may use forecasting, anomaly detection, optimization, rules, simulation, or retrieval-supported reasoning, but it must produce more than a suggestion. It must show the evidence used, confidence level, assumptions, alternatives considered, expected operational effect, financial consequence, and compliance constraints. A recommendation that cannot be explained, challenged, or audited should not be operationalized in a safety-critical maintenance environment.


The fourth layer is human authority and compliance governance. AI may accelerate diagnosis, but it must not displace certified responsibility, quality authority, or approved release-to-service accountability. This layer defines who may approve, reject, override, escalate, or suspend a recommendation. It also records why a recommendation was accepted or rejected. The governance issue is not whether humans are retained only ceremonially; it is whether the right accountable professional participates in the right decision cycle at the right level of risk.

The fifth layer is KPI-based feedback. Each decision should return evidence to the system: whether the bottleneck diagnosis was correct, whether the action reduced delay, whether a compliance boundary was activated, whether margin improved or deteriorated, whether the customer commitment was protected, and whether the AI recommendation was accepted, modified, or overridden. This turns the architecture into a learning decision system rather than a static reporting mechanism.

5. KPI and Control Logic

AMRO-DIGF treats KPIs as governance controls rather than decorative indicators. Turnaround risk can be defined as the probability that a maintenance event or repair order will exceed the committed completion time, weighted by operational criticality. Bottleneck severity can be expressed as the delay contribution of a constraint divided by total delay exposure. Parts readiness can be measured as the proportion of required parts available, serviceable, traceable, and allocated before task execution. Margin leakage can be estimated as actual job margin minus planned job margin after rework, overtime, expedite fees, exchange cost, warranty exposure, and penalty effects. AI decision quality can be measured through recommendation acceptance, override frequency, explanation completeness, post-decision accuracy, and compliance incident rate.

These KPI constructs should be connected to decision thresholds. For example, a high turnaround-risk score should trigger cross-functional review if parts readiness is low and supplier lead-time volatility is high. A margin-leakage alert should trigger commercial review if scope creep, rework, expedite shipping, or warranty uncertainty crosses predefined tolerance. A compliance-risk alert should block operational execution until authorized review



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



is complete. The key design principle is that KPI deviation should not merely appear on a dashboard; it should activate a governed decision pathway.

Control logic must also preserve proportionality. Low-risk recommendations may be approved by a planner or supervisor under predefined rules. Medium-risk recommendations require quality, engineering, logistics, or finance review depending on the constraint. High-risk recommendations affecting airworthiness, release-to-service, regulatory records, or safety-critical substitution require explicit authorized approval. This proportional decision-rights structure protects both agility and accountability.

6. Implementation and Governance

Implementation should begin with a narrow, high-value maintenance decision pathway rather than with a broad AI transformation program. Suitable first pathways include parts-caused turnaround delay, quality-hold resolution, supplier-delay escalation, repair-versus-replace recommendation, and negative-margin repair detection. Each pathway should be mapped across evidence sources, decision owners, approval gates, KPI thresholds, escalation rules, and feedback records.


The recommended roadmap has four waves. Wave 1 establishes data readiness: master-data cleanup, work-order lineage, part-number governance, supplier event capture, documentation completeness, and data-quality service levels. Wave 2 builds diagnostic intelligence: bottleneck taxonomy, delay attribution, exception rules, and root-cause coding. Wave 3 introduces recommendation intelligence: forecast-assisted prioritization, scenario comparison, constrained optimization, and explainable decision packs. Wave 4 institutionalizes governance: human authorization, override logging, compliance gates, AI risk review, cybersecurity monitoring, and post-decision learning.

RACI logic is essential. Accountable maintenance managers should own operational decisions; quality and continuing-airworthiness authorities should own compliance boundaries; logistics leaders should own parts readiness and supplier escalation; finance should own margin-leakage controls; data governance should own lineage, quality, and access rules; and digital or AI teams should own model monitoring, explainability, and technical controls. This prevents AI governance from becoming an IT-only responsibility.

Cybersecurity and privacy controls are not peripheral. Decision intelligence depends on trustworthy data flows, role-based access, change control, secure integrations, audit logs, and resilience against manipulation of maintenance or supplier data. A compromised recommendation engine or corrupted part-status record can create operational and safety consequences. Therefore, AMRO-DIGF should be aligned with enterprise information-security governance as well as AI risk management.

7. Discussion

The theoretical contribution of AMRO-DIGF is that it reframes aviation maintenance AI as decision governance. Much of the maintenance AI literature focuses on prediction quality, fault detection, or optimization. Those are necessary but incomplete. In safety-critical



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



کد اختصاصی نمایه ISC
۰۴۲۵۰-۲۲۷۹۹



maintenance, the relevant performance question is not only whether a model predicts correctly, but whether the organization can convert prediction into authorized, traceable, compliant, financially disciplined action. This shifts the analytical center from model output to decision episode.

The practical contribution is a route from ERP to decision intelligence. ERP remains indispensable because it preserves the transactional and compliance record. However, ERP becomes strategically more valuable when it is connected to diagnostics, recommendations, decision rights, and KPI feedback. This article therefore does not replace ERP with AI; it proposes an architecture in which ERP evidence becomes governable decision intelligence.

The relationship with prior AISA-L work is direct but differentiated. AISA-L established the logic of real-time KPI orchestration in airline logistics. AMRO-DIGF specializes that logic for aviation maintenance, where airworthiness, release authorization, quality control, part traceability, and margin discipline create stricter governance boundaries. This specialization is the article's main industry value: it translates a broad agentic logistics architecture into a maintenance-specific governance system that can be evaluated and implemented incrementally.

The main limitation is the absence of live implementation data. This is appropriate for a design-science conceptual article but must be transparent. Future research should test AMRO-DIGF through retrospective work-order data, anonymized repair-case evidence, simulation, and expert-panel validation. The most valuable empirical design would compare baseline ERP reporting with AMRO-DIGF-supported decision episodes across turnaround time, delay attribution accuracy, parts readiness, approval latency, compliance events, margin leakage, and override quality.

8. Conclusion

This article developed AMRO-DIGF as a KPI-governed architecture for auditable AI decision intelligence in aviation maintenance, repair and overhaul. The practical problem is that maintenance organizations possess extensive ERP, inventory, quality, finance, and operational data, but often lack a governed mechanism for converting evidence into timely, compliant, financially disciplined decisions. AMRO-DIGF integrates data lineage, diagnostics, AI-supported recommendation, human authority, compliance governance, and KPI feedback into a closed decision cycle.

The central conclusion is conservative and operational: AI creates credible value in aviation maintenance only when it improves decision quality without weakening airworthiness accountability. A recommendation that cannot show its evidence, assumptions, confidence, constraint checks, owner, approval route, and outcome is not mature decision intelligence. AMRO-DIGF therefore treats AI as a governed co-decision support layer, not as a substitute for authorized professional judgment, quality assurance, or accountable management.

For practice, the starting point should be a bounded decision pathway such as parts-caused turnaround delay, quality-hold resolution, supplier-delay escalation, repair-versus-replace decision-making, or negative-margin repair detection. For research, the next step is empirical evaluation through digital-twin simulation, expert validation, and longitudinal case

comparison. The ultimate objective is not merely faster maintenance; it is safer, more accountable, financially disciplined, and evidence-governed maintenance decision execution.

References

- [1] International Air Transport Association. Digital Aircraft Operations. Montreal: IATA; 2026. Available from: <https://www.iata.org/en/programs/ops-infra/techops/digital-aircraft-operations/>
- [2] European Union Aviation Safety Agency. Artificial Intelligence Roadmap 2.0: A human-centric approach to AI in aviation. Cologne: EASA; 2023.
- [3] Federal Aviation Administration. 14 CFR Part 145: Repair Stations. Washington (DC): Federal Aviation Administration; 2026. Available from: <https://www.ecfr.gov/current/title-14/chapter-I/subchapter-H/part-145>
- [4] European Union Aviation Safety Agency. Easy Access Rules for Continuing Airworthiness, including Part-145 maintenance organisation approvals. Cologne: EASA; 2025.
- [5] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Gaithersburg (MD): NIST; 2023.
- [6] International Organization for Standardization. ISO/IEC 42001:2023 Information technology - Artificial intelligence - Management system. Geneva: ISO; 2023.
- [7] Component Control. Quantum Control aviation software for the aviation aftermarket, including MRO, parts and logistics, manufacturing, and aircraft services. San Diego (CA): Component Control; 2026. Available from: <https://www.componentcontrol.com/>
- [8] Hevner AR, March ST, Park J, Ram S. Design science in information systems research. *MIS Quarterly*; 28(1):75-105, 2004.
- [9] Peffers K, Tuunanen T, Rothenberger MA, Chatterjee S. A design science research methodology for information systems research. *Journal of Management Information Systems*; 24(3):45-77, 2007.
- [10] Gregor S, Hevner AR. Positioning and presenting design science research for maximum impact. *MIS Quarterly*; 37(2):337-355, 2013.
- [11] Pavlyuk D, Alomar I. Artificial intelligence technologies for aircraft maintenance: A systematic literature review. *International Journal of Prognostics and Health Management*; 17(1), 2026. doi:10.36001/ijphm.2026.v17i1.4567.
- [12] Kabashkin I, Perekrestov V, Pivovarov M. AI-driven fault detection and maintenance optimization for aviation technical support systems. *Processes*; 13(3):666, 2025. doi:10.3390/pr13030666.
- [13] Shukla A, Bansal N. AI prognostics for aircraft fleet management: A strategic framework for operational efficiency and ROI. *International Journal of Applied Mathematics*; 38(4s):1363-1373, 2025. doi:10.12732/ijam.v38i4s.694.
- [14] MoghadasNian S, NaserPour Asiabari M, HeidariYekta A. AISA-L: An agentic AI strategy architecture for real-time KPI orchestration in sustainable, resilient airline logistics. *SSRN*; 2025. doi:10.2139/ssrn.6025895.
- [15] MoghadasNian SAH, Kashian H. AI strategy agent for airline logistics: A multi-layered, KPI-governed architecture for real-time optimization and ethical orchestration. 2nd International Conference on Artificial Intelligence in the Age of Digital Transformation. Zenodo; 2025. doi:10.5281/zenodo.18184458.
- [16] MoghadasNian SAH, Sarvi N. Enhancing MRO efficiency and effectiveness in the airline industry: A strategic approach to KPI-driven management. 9th International Conference on



National Conference on Information Technology, Nanotechnology,
Artificial Intelligence and Technological Futures Studies

Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park



Interdisciplinary Studies in Management & Engineering. Zenodo; 2024.
doi:10.5281/zenodo.19674936.

[17] MoghadasNian SAH, MahMoudy M. Airline logistics AI performance framework: A 360-degree, multi-layered KPI approach for safety, sustainability, efficiency, and innovation. 20th Iranian International Industrial Engineering Conference. Zenodo; 2025. doi:10.5281/zenodo.18251742.

[18] MoghadasNian SAH, MahMoudy M. Stock control in airline logistics: AI-driven inventory optimization for spare parts. 3rd International Conference on Recent Advances in Engineering, Innovation and Technology. Zenodo; 2025. doi:10.5281/zenodo.18372296.

[19] MoghadasNian SAH. AI-driven inventory optimization in airline logistics: Enhancing efficiency, sustainability, and operational performance. SSRN; 2025. doi:10.2139/ssrn.6119086.

[20] MoghadasNian SAH. AI-powered predictive maintenance in aviation operations. SSRN; 2025. doi:10.2139/ssrn.6117408.

[21] MoghadasNian SAH. Digital transformation in airline logistics: Enhancing operational efficiency through AI-driven predictive analytics and blockchain integration. SSRN; 2025. doi:10.2139/ssrn.6129506.

[22] MoghadasNian S, Rajol M, HosseinZadehShirazi Z. Strategic transformation through data governance: Leveraging Key Performance Indicators (KPIs) in the airline industry. 2nd International Conference on Engineering and Information Technology. Zenodo; 2024. doi:10.5281/zenodo.18728867.

[23] MoghadasNian SAH, RoshanSourat A. Enhancing airline maintenance efficiency: A KPI-driven framework for production control. 8th National Conference on Interdisciplinary Research in Engineering and Management. Zenodo; 2024. doi:10.5281/zenodo.19949130.

[24] Myllyaho L, Raatikainen M, Mannisto T, Mikkonen T, Nurminen JK. Systematic literature review of validation methods for AI systems. Journal of Systems and Software; 181:111050, 2021. doi:10.1016/j.jss.2021.111050.

[25] Cummins L, Sommers A, Bakhtiari Ramezani S, Mittal S, Jabour J, Seale M, Rahimi S. Explainable predictive maintenance: A survey of current methods, challenges and opportunities. arXiv; 2024. doi:10.48550/arXiv.2401.07871.