National Conference on Information Technology, Nanotechnology, Artificial Intelligence and Technological Futures Studies
Ministry of Science, Research and Technology
East Azarbaijan Science and Technology Park
ISC کد اختصاصی نمایه
۰۴۲۵۰-۲۲۷۹۹
وزارت علوم تحقیقات و فناوری
پارک علم و فناوری آذربایجان شرقی

Aviation AI Conscience Governance 360: A Design-Science Framework for Accountable, Safety-Critical and Data-Governed Airline AI

SeyyedAbdolHojjat MoghadasNian*

1. Tarbiat Modares University, s14110213@gmail.com

Abstract

Artificial intelligence (AI) is moving from experimental analytics into operational airline decision pathways, including operations control, predictive maintenance, disruption recovery, crew planning, passenger communication, revenue management, safety analytics, procurement and maintenance, repair and overhaul (MRO) logistics. The governance problem is therefore no longer whether AI can support airline performance, but whether AI-enabled recommendations remain safe, lawful, explainable, data-qualified, ethically defensible, economically justified and human-accountable. This article develops Aviation AI Conscience Governance 360 (AICG-360) as a design-science framework for accountable airline AI. The framework treats conscience not as machine morality, but as an executive governance architecture linking evidence, data quality, decision authority, risk thresholds, legal obligations, ethical constraints, economic value and post-decision learning. Drawing on aviation AI guidance, risk-management standards, data-quality models, design-science research and KPI-driven airline governance literature, the paper proposes eight governance dimensions: safety conscience, data conscience, model conscience, human-authority conscience, legal-compliance conscience, ethical conscience, economic-value conscience and institutional-learning conscience. It further specifies decision gates, escalation triggers, KPI domains, accountability roles and dashboard logic for airline boards, accountable managers, safety leaders, CIOs, CDOs, operations executives and compliance officers. The article contributes by reframing airline AI governance as auditable decision execution rather than isolated compliance, model validation or productivity automation; by operationalizing data quality as the epistemic foundation of responsible AI; and by offering a publishable artifact for future Delphi validation, case-study testing and dashboard prototyping.

Key words: Aviation AI governance, AICG-360, Safety-critical AI, Data governance, Human-AI teaming, Decision auditability

1. Introduction

1.1 Background and context

Commercial aviation is a high-reliability, safety-critical and data-intensive socio-technical system in which operational decisions depend on the alignment of aircraft, crews, maintenance status, airport constraints, passenger flows, commercial commitments, regulatory obligations and real-time disruption signals. AI increasingly enters this system through recommendation engines, anomaly detection, predictive maintenance, irregular-operations recovery, crew pairing, dynamic pricing, digital identity, service personalization, fraud detection and safety analytics. Even where AI does not formally decide, it can shape human judgement by ranking options, filtering alerts, producing risk scores, prioritizing cases and framing the apparent best operational action.

The why-now driver is explicit: aviation AI adoption, aviation-specific AI guidance, risk-based regulation and airline data-dependency are converging. EASA's AI Roadmap 2.0 positions aviation AI within a human-centric safety and ethics agenda [1], while EASA's Concept Paper Issue 2 expands guidance for Level 1 and Level 2 machine-learning applications, including learning assurance, explainability, ethics-based assessment and human oversight [2]. The EU AI Act creates a risk-based legal regime for AI systems placed on or used in the European Union [3]. NIST's AI Risk Management Framework provides a socio-technical vocabulary for trustworthy AI risk management [4]. SITA's Air Transport IT Insights further demonstrates that airline AI scaling is inseparable from data integration, data consistency and enterprise IT modernization [5]. ICAO working-paper discussions similarly foreground the need to preserve human control, technical robustness, transparency, resilience and accountability in aviation AI adoption [6].

This article defines AI as computational systems capable of generating predictions, classifications, recommendations or content that influence airline decision processes. Airline AI governance refers to the structures, roles, controls, metrics and accountability mechanisms through which airlines ensure that AI systems remain safe, lawful, reliable, explainable, ethical, cyber-resilient, economically justified and aligned with operational authority. AI Conscience Governance, the central construct of this article, is an auditable decision architecture that continuously asks whether an AI-enabled airline recommendation is evidence-based, data-qualified, risk-bounded, human-authorized, legally compliant, ethically defensible and institutionally learnable. The term conscience is used as a governance metaphor for accountable human judgement and organizational memory; it does not imply that AI possesses moral agency.

The primary article lens is Design Science Research (DSR), because the paper constructs and justifies a governance artifact rather than estimating a causal effect [7-9]. The framework is also informed by socio-technical systems theory and high-reliability organizing because AI-enabled aviation decisions emerge from interactions between technology, data, human factors, operating routines, regulatory constraints and organizational incentives [10-12]. A rival technology-adoption or productivity lens can explain why airlines adopt AI, but it is insufficient to govern accountability transfer, safety-margin erosion, passenger-rights exposure, data-quality legitimacy and auditability under operational time pressure.

1.2 Statement of problem

Most airlines do not lack AI ambition, dashboards or digital-transformation narratives. The more serious problem is fragmentation. IT governance may focus on architecture, vendor access and cyber controls. Data governance may focus on data ownership, cataloguing, master data, lineage and quality. Safety management systems may focus on hazards, occurrences, mitigation and assurance. Compliance teams may focus on regulation, privacy, contracts and auditability. Business units may focus on punctuality, cost reduction, revenue, customer experience and productivity. AI teams may focus on model accuracy, validation, drift, explainability and deployment. Yet AI-enabled airline decisions cut across all these domains simultaneously.

For example, an AI-driven disruption-management recommendation may affect passenger rights, crew legality, aircraft rotation, slot protection, hotel cost, customer communication, social-media exposure, downstream maintenance planning and operational safety buffers. A predictive-maintenance model may influence AOG avoidance, spare-parts allocation, MEL/CDL risk, MRO planning, aircraft availability and maintenance cost. A passenger-prioritization model may improve revenue or service recovery while creating fairness, transparency or discrimination concerns. The operational problem is therefore not simply whether a model is accurate; it is whether the whole decision pathway remains governable.

The missing construct is a 360-degree airline AI conscience system: a governance architecture that links data quality, model assurance, human authority, legal compliance, ethical constraint, safety accountability, economic logic and organizational learning into a single decision-execution regime. Without such architecture, AI adoption may produce decision pollution: outputs that appear technically sophisticated but are built on incomplete data, weak lineage, untested assumptions, opaque model behaviour, automation bias or unclear accountability.

1.3 Research gap

Existing frameworks provide indispensable foundations. EASA contributes aviation-specific AI assurance and human-AI teaming guidance [1,2]. The EU AI Act contributes legal risk classification and deployer/provider obligations [3]. NIST AI RMF contributes general AI risk-management vocabulary [4]. ISO/IEC 23894 contributes AI risk-management guidance, ISO/IEC 42001 contributes an AI management-system architecture, and ISO/IEC 25012 contributes a structured data-quality model [14-16]. DAMA-DMBOK further supports enterprise data-management and data-governance discipline [17]. However, these frameworks are not yet integrated into an airline-specific executive decision architecture that converts principles into accountable decision gates, KPI thresholds, escalation logic and board-level oversight.

The decisive gap is therefore not the absence of AI principles. It is the weak conversion of principles into auditable airline decision execution. Current discourse often treats AI ethics, data quality, safety assurance, model validation, legal compliance and operational performance as adjacent domains. In aviation, they are inseparable. The theoretical gap concerns how socio-technical AI risk becomes governable in a safety-critical commercial airline context. The methodological gap concerns how a design-science artifact can integrate standards, regulatory expectations, KPI logic and operational decision rights. The practical gap concerns how airline executives can see, challenge, approve, override and learn from AI-enabled decisions before responsibility silently migrates from accountable humans to algorithmic systems.

1.4 Research purpose, questions and objectives

The purpose of this article is to design AICG-360 as a 360-degree, aviation-specific, KPI-governed and audit-ready framework for responsible AI-enabled airline decision-making. The framework is intended to preserve innovation while preventing the erosion of safety, accountability, data integrity, human authority, ethical legitimacy and regulatory compliance.

The primary research question is: How can airlines design a 360-degree AI conscience governance framework that enables AI-driven decision support while preserving safety, accountability, data integrity, human authority, ethical legitimacy and regulatory compliance?

The secondary research questions are: What governance gap exists between current airline AI adoption and the requirements of safety-critical, human-centric and accountable AI? Which governance dimensions are necessary for an aviation-specific AI conscience framework? How can data quality be operationalized as a moral-epistemic foundation for trustworthy airline AI? How can AI governance be structured as an evidence-authority-risk-precedent-accountability system? Which KPIs, thresholds and escalation triggers should appear in an executive Aviation AI Conscience Dashboard? How can AICG-360 be embedded into existing safety, data, compliance, digital-transformation and performance-governance structures?

The objectives are to synthesize aviation AI governance requirements, design the AICG-360 artifact, define governance dimensions and decision rights, operationalize KPI logic and propose an implementation roadmap for airline executives and future empirical validation.

1.5 Significance and contributions

The significance of this study lies in its attempt to convert responsible AI from a principles discourse into an airline governance artifact. For aviation management, the article reframes AI as a decision-governance object embedded in operations, safety, commercial strategy, finance, customer experience and compliance. For AI governance, it translates general risk frameworks into the specific realities of time-sensitive, safety-critical and multi-actor airline decision-making. For data governance, it positions data quality as the epistemic foundation of AI legitimacy rather than as a back-office data-management concern. For safety management, it provides a bridge between AI assurance, model drift, automation bias and Safety Management Systems. For executive governance, it proposes dashboard logic that makes AI risk visible at board and accountable-manager levels.

This paper contributes by (1) proposing AICG-360 as a checkable design-science artifact with eight defined governance dimensions, decision gates and KPI domains; (2) operationalizing AI conscience as an auditable human-accountability architecture rather than an anthropomorphic claim about machines; (3) specifying how airline AI decisions should be evaluated through evidence quality, data lineage, model behaviour, human authority, legal obligations, ethical constraints, economic value and institutional learning; and (4) providing a validation pathway through Delphi assessment, expert interviews, airline case studies and dashboard prototyping.

1.6 Scope and boundary conditions

The paper focuses on commercial airlines and airline-related aviation ecosystems. The unit of analysis is the AI-enabled airline decision pathway, not the isolated algorithm. The framework applies most directly to AI systems used in operations control centres (OCC), flight-operations support, predictive maintenance and MRO, disruption management, crew planning, revenue management, customer service, digital identity, safety analytics, procurement and spare-parts optimization. Data-source examples include PSS, DCS, GDS/NDC, AODB, OCC platforms, MRO/CAMO systems, ERP, CRM, finance systems, crew systems, safety-reporting systems and cyber-security logs.

The argument is expected to hold where AI materially influences airline operational, safety, customer, legal or financial decisions. It is less directly applicable to purely administrative automation with no material risk, or to fully certified avionics systems where aviation product-certification pathways dominate organizational governance. The paper does not claim causal evidence, completed field validation or universal generalizability. It proposes a design-science artifact for future empirical testing.

2. Literature Review and Theory

2.1 Review logic

This review is an integrative and theory-building literature review supporting a design-science artifact. It combines authoritative standards and regulatory sources with aviation AI guidance, safety-critical systems literature, human-AI teaming research, data-governance standards and KPI-driven aviation research. The review prioritizes sources directly relevant to aviation AI governance, AI risk management, airline data quality, decision accountability, safety assurance and KPI-based governance. It is not presented as a PRISMA systematic review because its purpose is not to estimate an effect size, but to synthesize mechanisms and governance requirements that justify the AICG-360 artifact.

Evidence strength is assessed qualitatively. Official regulatory and standards documents are treated as strong normative and compliance evidence; peer-reviewed design-science and socio-technical research is treated as strong theoretical evidence; industry reports are treated as useful adoption evidence but not as proof of causal performance improvement; and the author's prior KPI-driven airline work is used as bounded conceptual continuity rather than independent empirical validation.

2.2 Theoretical background: foundation and rivals

DSR provides the article's construction logic because the research aims to solve a relevant organizational problem through a purposeful artifact, supported by rigorous grounding and evaluation planning [7-9]. AICG-360 is therefore not a causal model claiming observed performance effects. It is a governance artifact intended for future expert validation, case-study testing and dashboard prototyping.

Socio-technical systems theory provides the organizational lens because airline AI decisions emerge from interactions among algorithms, data pipelines, human operators, safety culture, regulatory constraints, workflows, incentives and time pressure [10]. High-reliability organizing adds a safety-critical lens because aviation depends on anticipation, operational sensitivity, resilience and disciplined escalation under uncertainty [11,12]. AI can strengthen these capabilities by improving signal detection and scenario analysis, but it can also weaken them if it creates automation bias, false precision, de-skilling or opaque authority transfer. Human-automation research warns that automated systems can be used, misused, disused or abused depending on design, training, trust and organizational control [13].

A rival technology-acceptance lens explains willingness to use AI but is insufficient because adoption does not equal safe governance. A productivity or resource-based lens can explain why AI may improve efficiency and competitive advantage, but it does not adequately specify who remains accountable when AI recommendations affect safety margins, passenger rights or regulatory obligations. A compliance-only lens also fails because legal conformity does not automatically guarantee operational wisdom, ethical legitimacy or data reliability. The chosen lens therefore integrates governance, safety and decision accountability rather than treating AI as a standalone digital capability.

2.3 Conceptual domain map

AICG-360 is built around eight constructs. Safety conscience is the capacity to identify, assess and escalate AI-related safety implications. Data conscience is the capacity to verify that data are complete, current, traceable, representative, semantically consistent and operationally valid. Model conscience is the capacity to monitor model validity, robustness, explainability, drift, bias, cybersecurity exposure and lifecycle control. Human-authority conscience is the preservation of accountable human judgement, challenge rights, override quality and escalation discipline.

Legal-compliance conscience is the systematic mapping of AI use cases to relevant legal, regulatory, contractual, privacy and audit obligations. Ethical conscience concerns fairness, dignity, transparency, non-discrimination, stakeholder harm and proportionality beyond minimal legal compliance. Economic-value conscience requires AI benefits to be measured against assurance cost, safety buffers, compliance exposure, passenger impact and lifecycle risk. Institutional-learning conscience is the capacity to capture precedents, audit findings, incidents, overrides, near-misses and model lessons into future governance. These constructs operate at the organizational and decision-pathway level, not at the isolated algorithm level. A model can be technically strong but governance-weak if it lacks data lineage, human escalation, regulatory mapping or post-decision learning.

2.4 Critical synthesis by mechanisms

Aviation AI assurance and safety-critical governance. The strongest consensus in the reviewed material is that aviation AI must be human-centric, safety-oriented and assurance-driven. EASA's AI roadmap and concept paper establish a trajectory in which aviation AI requires learning assurance, explainability, ethics-based assessment and human oversight [1,2]. ICAO-related discussions similarly emphasize informed dialogue on AI's evolving role in aviation and the need to preserve human control, robustness, transparency, resilience and accountability [6]. The contested issue is not whether oversight is needed, but how operationally specific it should be. Generic guidance can require oversight without specifying how an OCC duty manager, CAMO leader, safety manager, revenue manager or accountable executive should challenge AI under time pressure. Implication for AICG-360: the framework must define decision gates, escalation thresholds and role-specific accountability rather than merely repeating principles.

Trustworthy AI, risk management and management-system integration. NIST AI RMF, ISO/IEC 23894 and ISO/IEC 42001 converge on the view that AI risk is lifecycle-based, socio-technical and organizational rather than purely technical [4,14,15]. The debate concerns whether existing management systems can absorb AI risk or whether AI requires a specialized governance layer. This paper resolves the debate by treating AICG-360 as an integrative overlay, not a replacement for Safety Management Systems, compliance, IT governance, data governance or enterprise risk management. Implication for AICG-360: the framework must provide crosswalk logic between existing airline governance systems and AI-specific risk controls.

Data quality as epistemic and ethical foundation. Airline AI depends on heterogeneous and often imperfect data from OCC, AODB, PSS, DCS, GDS/NDC, MRO/CAMO, ERP, finance, CRM, crew, safety and logistics systems. ISO/IEC 25012 defines a general data-quality model, while DAMA-DMBOK supports wider data-management discipline [16,17]. In aviation, the ethical problem often begins before the model: delayed maintenance data, inconsistent aircraft status, duplicated passenger records, incomplete service-recovery notes or misclassified safety events can distort AI output. Prior airline data-governance work similarly argues for KPI-driven governance, analytics discipline and data-quality accountability in airline management [19,23]. Implication for AICG-360: data conscience must be a hard governance gate, not a background IT hygiene metric.

Human-AI teaming, automation bias and authority preservation. Human-AI teaming literature emphasizes that AI systems should support human judgement rather than replace contextual accountability. In aviation, judgement often depends on tacit knowledge, safety margins, crew legality, passenger care, regulatory interpretation and disruption context. The risk is not only model error, but automation bias: the tendency to over-trust automated recommendations because they appear precise, ranked or technically authoritative [13]. Implication for AICG-360: human oversight must be measured by quality of challenge, override rationale and decision traceability, not merely by the formal presence of a human in the loop.

Economic value, KPI governance and responsible scaling. Airlines have strong incentives to scale AI because AI can improve prediction, speed, disruption recovery, asset utilization, personalization and cost control. Previous airline research connects AI, digital transformation, predictive maintenance, human-centric AI, MRO intelligence, logistics optimization and KPI-driven governance [19-28]. However, productivity gains are not sufficient evidence of responsible AI. A model that reduces delay cost but increases passenger-rights exposure, safety-buffer erosion, opaque discrimination or audit risk may be value-destructive at enterprise level. Implication for AICG-360: KPI governance must connect AI outputs to decision quality, not only operational efficiency.

2.5 Measurement and KPI operationalization

The central measurement challenge is to avoid reducing AI governance to a single maturity score. AICG-360 uses a KPI portfolio because airline AI decisions create multiple simultaneous consequences. Safety conscience can be measured through AI-related safety occurrences, safety-buffer exceptions, AI-influenced hazard reports, near-miss learning cycle time and percentage of AI use cases mapped to SMS risk assessment. The validity threat is under-reporting; mitigation requires confidential reporting, safety-culture integration and mandatory AI-decision tagging in relevant systems.

Data conscience can be measured through data completeness, freshness, lineage coverage, duplicate rate, master-data conformance, semantic consistency and representativeness of training and validation datasets. The validity threat is that system-level data quality may mask use-case-level defects; mitigation requires use-case-specific data-quality service-level agreements and lineage checks. Model conscience can be measured through model-validation pass rate, drift frequency, explainability sufficiency, bias-test completion, cybersecurity-control compliance, rollback readiness and lifecycle documentation completeness. The validity threat is that historical validation may not generalize to disruption regimes; mitigation requires scenario testing, stress testing and drift monitoring.

Human-authority conscience can be measured through documented override rate, quality of override rationale, challenge-right usage, escalation adherence, decision-right clarity and post-event accountability traceability. The validity threat is that low override rates may indicate either strong AI or weak human challenge; mitigation requires review of override quality and false-negative challenge cases. Economic-value conscience can be measured through net benefit after assurance cost, avoided delay cost, AOG avoidance, cost-per-decision improvement, service-recovery value, compliance exposure avoided and benefit-risk ratio. The validity threat is overstated benefit if externalities are excluded; mitigation requires safety, legal, reputational and customer-cost adjustments.

A possible AICG-360 score can be expressed as a weighted average of Safety, Data, Model, Human Authority, Legal Compliance, Ethical, Economic Value and Institutional Learning domain scores, with hard red gates for safety-critical or unlawful conditions. A high aggregate score cannot override a red safety, legal or ethical gate. This prevents compensatory scoring from hiding unacceptable risk.

3. Methodology

3.1 Research design and rationale

This article uses DSR because the objective is to create and justify a governance artifact rather than to test a causal hypothesis with primary field data. The artifact is AICG-360, a framework for airline AI decision governance. The claim type is design/evaluation rather than causal. The article claims that AICG-360 is a coherent, theoretically grounded and operationally relevant architecture for governing AI-enabled airline decisions. It does not claim that the framework has already improved safety, cost, compliance or passenger outcomes.

The DSR process followed five logic stages: problem diagnosis, knowledge-base synthesis, artifact construction, internal coherence evaluation and future validation planning. Problem diagnosis identified fragmentation across AI, safety, data, compliance, IT and business governance. Knowledge-base synthesis connected aviation AI guidance, trustworthy AI, data-quality standards, DSR, socio-technical theory, high-reliability organizing and KPI-driven airline research. Artifact construction specified the eight dimensions, decision gates, KPI domains, escalation triggers and ownership logic. Internal coherence evaluation checked whether each dimension links to the research gap and decision pathway. Future validation planning specifies Delphi, case-study, dashboard-pilot and operational-data testing.

3.2 Study context, unit of analysis and scope

The study context is commercial airline AI adoption in safety-critical, regulated and multi-system environments. The unit of analysis is the AI-enabled airline decision pathway, defined as the sequence from data capture and model output to human interpretation, authorization, execution, monitoring, audit and learning. The scope includes OCC, flight operations support, MRO/CAMO, disruption management, crew planning, revenue management, customer service, digital identity, safety analytics, procurement and logistics. Out of scope are certified flight-control automation, avionics product certification, causal airline performance modelling and confidential airline operational-data analysis.

3.3 Data sources, provenance and selection logic

The manuscript is based on a documentary and conceptual evidence base rather than primary fieldwork. Sources include aviation AI guidance, AI governance frameworks, data-quality standards, DSR and socio-technical literature, and selected airline KPI and AI-governance works listed in the reference section. No passenger records, employee records, confidential airline logs, operational datasets, model files or expert panel results are reported. The source hierarchy prioritizes official aviation and standards sources for normative requirements, peer-reviewed or canonical research for theoretical grounding, and SSRN or DOI-bearing author works for continuity with the KPI-driven airline research stream.

3.4 Artifact design and operationalization

Artifact construction followed a mechanism-to-governance translation logic. First, each governance requirement was mapped to a decision-risk question: Is the AI use case safety-relevant? Is the data decision-worthy? Is the model valid under current conditions? Is the human decision right clear? Is the legal obligation mapped? Is the ethical risk acceptable? Is the business value risk-adjusted? Is the decision learnable? Second, each question was translated into a governance dimension. Third, each dimension was translated into candidate KPIs, escalation triggers and accountable roles. Fourth, all dimensions were integrated into an executive dashboard and implementation roadmap.

The framework does not assume that all AI use cases require the same level of assurance. A chatbot handling generic information does not require the same escalation as a predictive-maintenance model influencing aircraft availability or a disruption-recovery model affecting passenger rights and crew legality. AICG-360 therefore applies proportional governance, with criticality, regulatory exposure, data maturity and time pressure determining gate intensity.

3.5 Validity, reliability and trustworthiness

Construct validity is addressed by mapping each AICG-360 dimension to recognized governance, safety, data, risk or accountability requirements. Internal validity is addressed through logic consistency between problem, gap, artifact, KPI domains and implementation pathways. External validity is not claimed; it requires future application across different airline sizes, regions, maturity levels and use-case classes. Reliability is supported by explicit definitions, decision-gate logic and KPI operationalization, but inter-rater reliability remains a future empirical requirement. Auditability is supported by traceable decision records, evidence gates, accountable owners and post-decision learning loops.

Because no human participants, operational datasets or confidential records are used in this conceptual article, ethics approval is not claimed. Future empirical work involving interviews, surveys, passenger data, employee records or operational logs would require informed consent, confidentiality controls, data minimization, access governance, GDPR alignment where applicable, and safety/regulatory review.

4. AICG-360 Framework and Findings

4.1 Framework logic

The central design principle of AICG-360 is: AI may recommend; accountable humans and institutions must authorize, monitor and learn. The framework rejects both extremes of AI rejection and uncritical AI delegation. It accepts AI as a decision-support capability while requiring that the decision pathway remains auditable, risk-bounded and human-accountable.

The framework contains eight dimensions. Safety conscience asks whether the AI use case protects or weakens aviation safety. Data conscience asks whether the evidence base is complete, timely, representative, traceable and fit for decision use. Model conscience asks whether the model is valid, monitored, explainable, robust, bias-tested and rollback-ready. Human-authority conscience asks whether decision rights, challenge rights, override protocols and accountability chains are explicit. Legal-compliance conscience asks whether applicable aviation, AI, privacy, passenger-rights, employment, procurement and contractual obligations are mapped. Ethical conscience asks whether AI use respects fairness, dignity, non-discrimination, transparency and stakeholder harm prevention. Economic-value conscience asks whether value is risk-adjusted and assurance-aware. Institutional-learning conscience asks whether incidents, overrides, audit findings and near-misses correct future data, models, policies and training.

4.2 Decision-gate architecture

AICG-360 organizes airline AI governance through sequential decision gates. Gate 1 is use-case criticality classification: classify the AI use case as low, medium, high or safety/rights critical. Gate 2 is data readiness: verify data lineage, completeness, freshness, semantic consistency and owner approval. Gate 3 is model assurance: verify validation, drift controls, explainability sufficiency, bias testing, cybersecurity and rollback. Gate 4 is human authority: define who may accept, challenge, override or escalate the recommendation. Gate 5 is legal and ethical review: verify applicable obligations and harm-prevention controls. Gate 6 is economic value: assess risk-adjusted benefit after assurance cost and externalities. Gate 7 is execution monitoring: tag AI-influenced decisions and monitor outcomes. Gate 8 is learning: feed incidents, overrides, audit findings and lessons into future governance.

This architecture is intentionally non-compensatory. A strong economic case cannot compensate for a red safety gate. High model accuracy cannot compensate for unlawful processing, weak data lineage or absent human authority. This makes the framework suitable for safety-critical and regulated environments where a single unacceptable control failure may invalidate the decision pathway.

4.3 Executive dashboard logic

The Aviation AI Conscience Dashboard should not be another productivity dashboard. It should show whether AI-enabled decision pathways are governable. Suggested board-level indicators include percentage of AI use cases classified by criticality; proportion of AI use cases with approved data lineage; model-drift exceptions; explainability sufficiency by use case; AI-influenced safety reports; human override rate and override-rationale quality; regulatory mapping completion; unresolved ethical-risk cases; risk-adjusted value realization; and corrective-action closure time.

Dashboard cadence should be risk-based. Safety-critical and high-operational-impact AI use cases require near-real-time or weekly monitoring by operational and safety owners, with monthly executive review. Medium-risk use cases may be reviewed monthly, and low-risk use cases quarterly. Board-level review should focus on risk concentration, unresolved red gates, overdue corrective actions, systemic data-quality weaknesses, model drift and AI-related incident learning.

4.4 Governance ownership and RACI logic

AICG-360 requires clear decision rights. The board is accountable for AI risk appetite, oversight and assurance expectations. The accountable manager is responsible for ensuring that AI use does not erode safety or regulatory accountability. The CIO owns architecture, integration and cyber controls. The CDO owns data governance, lineage, master data and data-quality SLAs. The safety manager owns SMS integration, AI-related hazard review and safety escalation. The compliance/legal function owns regulatory and contractual mapping. Business owners own use-case value, operational feasibility and process adoption. AI/model owners own model validation, monitoring, drift controls and documentation. Internal audit reviews whether the governance evidence exists and whether decisions can be reconstructed.

This RACI logic prevents responsibility laundering. AI outputs must not become organizational excuses. Every material AI-enabled decision pathway should have a named business owner, data owner, model owner, safety/compliance reviewer where applicable, escalation authority and audit trail.

4.5 Implementation roadmap

Implementation can proceed in four waves. Wave 1 establishes inventory and risk classification: identify AI use cases, decision pathways, data sources, model owners and affected stakeholders. Wave 2 builds minimum governance controls: data lineage, KPI definitions, model documentation, override protocols, compliance mapping and escalation routes. Wave 3 operationalizes the dashboard: integrate KPI feeds, risk thresholds, red-gate reporting and corrective-action tracking. Wave 4 institutionalizes learning: run post-decision reviews, calibrate thresholds, update training, refine model governance and embed the framework into SMS, data governance, enterprise risk management and internal audit.

A pragmatic starting point is to pilot AICG-360 in two contrasting use cases: one operationally critical use case such as predictive maintenance, disruption recovery or OCC decision support; and one passenger-facing use case such as personalization, service recovery or digital identity. This contrast tests whether the framework can handle both safety-operational and fairness/customer-rights risks.

5. Discussion

5.1 Interpretation against theory

AICG-360 extends DSR by translating a complex aviation governance problem into a structured artifact with dimensions, gates, KPIs, roles and validation pathways. It also operationalizes socio-technical theory by showing that responsible airline AI is not only a property of models, but of the wider decision ecology linking data, humans, workflows, organizations and regulators. From a high-reliability perspective, the framework reinforces preoccupation with failure, sensitivity to operations and disciplined escalation by making AI-related risk visible before and after execution.

The framework refines AI governance by shifting attention from isolated assurance to auditable decision execution. A technically validated model can still be unsafe or ethically weak if it is deployed with poor data, weak human authority, unclear legal mapping or no learning loop. Conversely, a modest model may be responsibly used if it is bounded, explainable, human-authorized, monitored and learnable. This distinction is central to aviation because responsibility cannot be delegated to the apparent sophistication of an algorithm.

5.2 Rival explanations and risks

A possible rival explanation is that existing Safety Management Systems, IT governance, compliance and data governance are already sufficient. AICG-360 does not reject those systems; it argues that AI-enabled decision pathways cross them simultaneously and therefore need an integrative overlay. Another rival explanation is that model accuracy is the primary governance concern. The article rejects this narrow view because model accuracy does not guarantee data legitimacy, legal compliance, fairness, human challenge, safety escalation or auditability.

The framework also carries risks. KPI governance may produce gaming if managers optimize dashboard scores rather than decision quality. Red-gate logic may slow innovation if thresholds are poorly calibrated. Human-in-the-loop governance may become symbolic if human reviewers lack time, authority, competence or psychological safety to challenge AI.

Data-quality scoring may hide local defects if aggregated at too high a level. These risks require governance culture, audit discipline and periodic threshold review.

5.3 Theoretical, methodological and practical implications

The theoretical implication is that airline AI governance should be understood as a socio-technical decision architecture, not as a generic adoption problem. The methodological implication is that DSR can generate sector-specific governance artifacts where causal data are not yet available but practical risk is urgent. The measurement implication is that AI governance should be assessed through a portfolio of safety, data, model, human authority, compliance, ethics, value and learning indicators rather than a single maturity score.

For practitioners, AICG-360 provides a board-to-frontline logic for responsible AI scaling. Boards can set risk appetite and red-gate principles. Accountable managers can ensure AI does not dilute safety accountability. CIOs and CDOs can align system architecture, data lineage and data-quality SLAs. Safety, compliance and legal leaders can embed AI into existing assurance and regulatory processes. Operations and commercial leaders can evaluate value only after safety, data, legal and ethical conditions are satisfied. Internal audit can test whether AI-influenced decisions can be reconstructed.

5.4 Responsible use

Responsible use of AICG-360 requires restraint. The framework should not be used to create a false sense of certainty, to justify cost reduction at the expense of safety buffers, or to convert ethical judgement into mechanical scoring. It should also not be used to claim regulatory compliance without legal review or empirical validation. Its most valuable function is to make the decision pathway visible: what data were used, what model produced the recommendation, who had authority, what risks were known, what obligations applied, what value was expected, what happened after execution and what was learned.

6. Conclusion

This article asked how airlines can design a 360-degree AI conscience governance framework that enables AI-driven decision support while preserving safety, accountability, data integrity, human authority, ethical legitimacy and regulatory compliance. The answer is AICG-360: a design-science governance artifact that integrates eight dimensions - safety, data, model, human authority, legal compliance, ethics, economic value and institutional learning - into an auditable decision-execution architecture.

The core contribution is not another AI ethics checklist. It is a sector-specific architecture for governing airline AI where decisions are time-sensitive, safety-critical, multi-system, legally exposed and commercially consequential. The framework clarifies that AI recommendation does not equal organizational authorization, model accuracy does not equal decision legitimacy, and economic value does not override safety, legal or ethical gates.

The paper remains conceptual and requires future validation through Delphi panels, expert interviews, airline case studies, dashboard pilots and operational decision-log analysis. Within that boundary, it provides a rigorous basis for airline executives, regulators and researchers seeking to scale AI without diluting human and institutional responsibility. This paper changes airline AI governance from a principles-and-compliance discussion into an auditable decision-execution literature and enables accountable, data-governed and safety-critical AI oversight in airline practice and policy.

References

[1] European Union Aviation Safety Agency. Artificial Intelligence Roadmap 2.0: A human-centric approach to AI in aviation. Cologne: European Union Aviation Safety Agency; 2023.

[2] European Union Aviation Safety Agency. Artificial Intelligence Concept Paper Issue 2: Guidance for Level 1 and Level 2 machine-learning applications. Cologne: European Union Aviation Safety Agency; 2024.

[3] European Parliament and Council of the European Union. Artificial Intelligence Act. Brussels: European Union; 2024.

[4] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg (MD): U.S. Department of Commerce; 2023.

[5] SITA. Air Transport IT Insights 2025. Geneva: SITA; 2025.

[6] International Civil Aviation Organization. Artificial intelligence contribution to aviation. ICAO Assembly, 42nd Session, Working Paper. Montreal: International Civil Aviation Organization; 2025.

[7] Hevner AR, March ST, Park J, Ram S. Design science in information systems research. MIS Quarterly; 28(1): 75-105, 2004.

[8] 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.

[9] Gregor S, Hevner AR. Positioning and presenting design science research for maximum impact. MIS Quarterly; 37(2): 337-355, 2013.

[10] Trist EL, Bamforth KW. Some social and psychological consequences of the Longwall method of coal-getting. Human Relations; 4(1): 3-38, 1951.

[11] Weick KE, Sutcliffe KM. Managing the unexpected: Sustained performance in a complex world (3rd ed.). San Francisco: Jossey-Bass; 2015.

[12] Reason J. Managing the risks of organizational accidents. Aldershot: Ashgate; 1997.

[13] Parasuraman R, Riley V. Humans and automation: Use, misuse, disuse, abuse. Human Factors; 39(2): 230-253, 1997.

[14] International Organization for Standardization. ISO/IEC 23894:2023 Artificial intelligence - Guidance on risk management. Geneva: ISO; 2023.

[15] International Organization for Standardization. ISO/IEC 42001:2023 Artificial intelligence - Management system. Geneva: ISO; 2023.

[16] International Organization for Standardization. ISO/IEC 25012:2008 Software engineering - Software product quality requirements and evaluation - Data quality model. Geneva: ISO; 2008.

[17] DAMA International. DAMA-DMBOK: Data management body of knowledge (2nd ed.). Basking Ridge (NJ): Technics Publications; 2017.

[18] International Air Transport Association. Exploring artificial intelligence and digital identity use cases in aviation. Montreal/Geneva: International Air Transport Association; 2025.

[19] MoghadasNian SAH. Digital transformation and AI in airline management: Aligning agility, innovation, and data-driven strategies. SSRN; 2025. https://doi.org/10.2139/ssrn.6137987

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

[21] MoghadasNian SAH. Human-centric AI integration in airline operations: Enhancing safety, efficiency, and workforce resilience. SSRN; 2025. https://doi.org/10.2139/ssrn.6099266

[22] MoghadasNian SAH. Generative AI for aviation maintenance: Optimizing predictive analytics through strategic prompt engineering and hybrid human-AI collaboration. SSRN; 2025. https://doi.org/10.2139/ssrn.6072392

[23] MoghadasNian SAH, Manafi F. Elevating airline performance with data analytics: A strategic guide to key performance indicators. SSRN; 2024. https://doi.org/10.2139/ssrn.6236539

[24] MoghadasNian S, Naserpour Asiabari M, Heidari Yekta A. AISA-L: An agentic AI strategy architecture for real-time KPI orchestration in sustainable, resilient airline logistics. SSRN; 2025. https://doi.org/10.2139/ssrn.6025895

[25] Moghadasnian SAH, Kashian H. Agentic AI in Airline Management: A KPI-governed architecture for trust-based autonomy, strategic co-leadership, and operational excellence. SSRN; 2025. https://doi.org/10.2139/ssrn.6033194

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

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

[28] MoghadasNian SAH, GhajarGar F. AI-powered passenger experience in airlines: A KPI-driven framework with mixed-methods evidence. SSRN; 2025. https://doi.org/10.2139/ssrn.6017454