SAFETY INTELLIGENCE IN AVIATION: THE NEW REALITY FOR PRESCRIPTIVE SAFETY MANAGEMENT

ALG

May 2026

INTRODUCTION

Aviation today stands among the safest modes of transport in the history of human mobility. Continuous improvements in safety management systems, rigorous regulation, and a deep safety culture across airlines, airports, and Air Navigation Service Providers (ANSPs) have driven accident and serious incident rates to historically low levels, even as traffic volumes continue to grow year on year. In 2024 the global commercial air transport system saw traffic surpass pre-pandemic levels, reflecting both the resilience of the industry and its commitment to safety performance despite increased demand on infrastructure and operations[1].

Modern aircraft and airport systems produce continuous streams of operational intelligence, from Flight Data Analysis programs, air traffic control, weather services, safety reports, and maintenance records. All together offer an unprecedented window into the 'safety health' of the aviation system. In practice, most of this data remains siloed within individual stakeholders, processed for compliance rather than insight, and analyzed retrospectively rather than in ways that could inform decisions before events occur. The gap between data availability and actionable intelligence is wide, and it has a measurable cost: in operational efficiency, in safety margins eroded by risks that were visible but unseen, and in the ability of operators, and their partners to manage risk with the precision that modern tools now make possible.

Closing that gap is what Safety Intelligence is about. Formally defined by the International Civil Aviation Authority (ICAO) as the outcome of systematically collecting, analyzing, and interpreting safety data to support data-driven decision making, Safety Intelligence represents a step change in how aviation stakeholders can use the information they already hold to move from reactive safety management toward proactive, systemic risk understanding.

This paper explores Safety Intelligence as a new paradigm in aviation safety management: what it means in practice, why it matters now, and how organizations can build toward it in a realistic, structured way.

[1]IATA, "IATA Annual Safety Report (2024) – Executive Summary"

AVIATION SAFETY: EXCELLENCE UNDER PRESSURE

Why a safe industry now faces new forms of complexity and systemic risk

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

The evolution of aviation safety

Aviation's safety has been built through a progressive re-framing of the problem, moving to new perspectives when the previous one reached its limits. In the mid-twentieth century, the focus was on hardware: more redundant systems, better materials, stricter certification. When reliability improvements stopped making a big difference, attention shifted to the human element such as Crew Resource Management, non-technical skills and the slow cultural dismantling of cockpit hierarchy as a safety liability. Later, the focus expanded even more: safety not as a property of aircraft or pilots, but as an emergent quality of systems, processes, and institutions.

The results speak for themselves. Collaborative risk reduction programs like the Commercial Aviation Safety Team demonstrated what systematic, data-informed intervention could produce at scale: an 83% reduction in the U.S. commercial aviation fatality rate between 1998 and 2007[2]. ICAO's formalization of Safety Management Systems in 2006 embedded these principles globally, making structured hazard identification and organizational learning a baseline expectation rather than a differentiator. By 2024, accident rates had fallen to 1.13 per million sectors, down from 3.72 in 2005[3], which is an outstanding achievement.

The evolution of aviation safety practices reflects a broader shift from reactive accident investigation toward more proactive and performance-oriented frameworks. In early phases, safety improvements were largely driven by technical reliability enhancements and corrective actions following accidents. Over time approaches expanded to incorporate systematic safety risk management, hazard identification, and organizational learning, culminating in the widespread adoption of Safety Management Systems (SMS) following formal guidance from ICAO in 2006[4]. These SMS frameworks embed structured policies, risk assessment, and assurance processes that encourage hazard identification before accidents occur[5].

But the logic of each previous transition applies here too. At some point, the existing framework delivers diminishing returns. The question is whether the industry recognizes that moment before or after the next serious wake-up call.

[2] Commercial Aviation Safety Team (CAST), "About CAST," accessed 2025

[3] IATA 2024 Annual Safety Report

[4] Al-Dmour, H., Alkhawaldeh, H., Al-Dmour, A. et al. "The integrated role of Safety Management Systems (SMS) and risk management in achieving aviation sustainability". Discov Sustain 6, 985 (2025).

[5] International Civil Aviation Organization (ICAO), Safety Management Manual (SMM), 4th ed., Doc 9859 (Montreal: ICAO, 2018).

A system under pressure

The operational environment that the aviation industry navigates nowadays is different from the one SMS frameworks were designed for. Traffic continues to grow, expected to double 2019 levels by 2045[6]. Turnaround times are tighter. Infrastructure, particularly in major hub airports, is running close to its design limits. Fleet diversification is accelerating, with increasing automation and the gradual integration of unmanned systems into an airspace that was never designed with them in mind. All these pressures over the system introduce new hazards that increases the risk in operations.

None of this is unmanageable on its own. The harder problem is that these pressures and their associated risks compound each other, and they do so in ways that existing monitoring systems were not built to detect.

On top of that, there is also a structural shift in what “safety risk” even means. Traditionally, the boundaries between safety and other risks (environmental, security, financial), were managed largely in isolation. The concept of Integrated Risk Management acknowledges that a single hazard can have consequences across multiple domains simultaneously, each requiring its own management response. It is not simply a matter of having a broader view, the risk itself has multiple active dimensions, and an effective response must address each axis without losing sight of how they interact.

A clear example are cybersecurity events. Not every cyberattack is a safety event. As an example, a breach that compromises customer payment data has no ‘safety impact’ on flight operations. However, an attack targeting ATM systems, aircraft data links or dispatch and crew coordination platforms entail safety hazards. GNSS jamming and spoofing, which are increasingly documented in conflict-adjacent airspace across multiple regions, degrades the navigation and timing infrastructure that modern operations depend on. These are recurring features of today’s operational landscape, and they require safety managers to hold a wider risk picture than flight data and occurrence reports alone can provide.

ICAO’s 2026–2028 Global Aviation Safety Plan explicitly addresses GNSS interference and cybersecurity resilience as systemic safety priorities. The Muscat Declaration, adopted in December 2024, called on states worldwide to treat cybersecurity with the same priority as other aspects of aviation safety. Regional frameworks such as EASA’s European Plan for Aviation Safety go further in operational detail, but the underlying shift in how risk is understood is a global one. The next paradigm shift in aviation safety is not primarily technical: it is about information, who has it, whether it flows to the right places, and whether it arrives in time to be useful in a system under pressure.

[6]ACI World Airport Traffic Forecasts 2025-2054

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

THE LIMITS OF TODAY’S SAFETY MANAGEMENT

Despite aviation’s access to rich operational data, fragmentation and reactive practices limit system-wide insight

Lots of data, little insight

Modern aviation generates vast volumes of operational data, far beyond what legacy systems were designed to manage. Flight Data Monitoring (FDM) programs, maintenance logs, air traffic control (ATC) data, safety reporting, weather feeds, and airport systems together produce terabytes of raw information daily. For example, integrated system monitoring across radar and weather sources alone can produce multiple terabytes of data per year even for a single airspace domain[7]. This data richness has the potential to improve safety performance and operational efficiency significantly, but it also creates serious challenges in data management, processing, and interpretation.

Despite the sheer scale of data, insight remains elusive. Much aviation data collection is driven by reporting obligations and compliance requirements, which ensure minimum standards but do not inherently produce actionable insight. Organizations often focus on generating and submitting the required information, rather than investing in the additional effort, analytical tooling, and cross-domain integration required to derive deeper understanding beyond compliance. Effective insight depends on the combination of data ingestion, rigorous processing, and advanced analytics, yet many operators (particularly those with limited budgets or technical capacity) struggle to make that transition.

One of the most persistent barriers to generating insight is the siloed nature of aviation safety data. Airlines, airports, ANSPs, and manufacturers each collect rich and relevant data, but these rarely flow beyond individual domains. Differences in system formats, lack of common standards, contractual and competitive concerns, and cultural reluctance to share sensitive information all contribute to fragmented data ecosystems[8]. Such silos impede a holistic understanding of safety events: an unstable approach incident, for example, can involve pilot actions, airline procedures, runway configuration, and ATC sequencing, yet no single party typically has visibility across all contributing domains. In the absence of integrated analysis, risk signals that span domains remain hidden.

This gap between data collection and insight generation is widely recognized. Data quality issues, inconsistent formats, and legacy systems complicate integration, while traditional analytical capacity often lags the volume and variety of incoming data streams. Without strong cross-stakeholder collaboration and modern analytical infrastructure, stakeholders are left with disconnected snapshots rather than a unified picture of safety risk, limiting the potential for proactive risk management and systemic performance improvement.

[7]NASA, In-Time Aviation Safety Management Systems: Challenges and Research for an Evolving Aviation System (Washington, DC: National Academies Press, 2023)

[8]"The Value of Data Sharing in the Future of Aviation," Albatros HF, accessed 2026.

Reactive safety practices

Despite the maturity of Safety Management Systems across the aviation industry, hazard identification and safety monitoring remain largely reactive and manual. In many organizations, safety risk is still primarily identified through mandatory and voluntary occurrence reporting, post-event investigations, and exceedance-based monitoring of predefined thresholds. These mechanisms are effective at capturing what has already happened[9]. As a result, safety analysis is often anchored to discrete occurrences rather than to evolving system behavior. The problem is that "what has already happened" is the wrong unit of analysis if the goal is prevention.

Rare but high-impact events (such as runway excursions or loss-of-control scenarios) may only materialize after a long accumulation of subtle contributing factors across operations, infrastructure, human performance, and environmental conditions. Traditional models, which rely on siloed datasets and event-based thresholds, struggle to capture these interactions or to identify emerging risk patterns before they manifest operationally[10]. The limited use of predictive analytics and cross-stakeholder models further constrains the ability to move from descriptive safety monitoring toward anticipatory risk management.

Evidence-Based Training (EBT) illustrates both progress and remaining limitations in current safety practices. EBT represents a significant evolution by using operational data to define pilot competencies and prioritize training scenarios based on observed risks rather than historical accident categories[11]. However, EBT remains reactive in nature, as the scenarios and competencies emphasized are still largely derived from events and deviations that have already occurred. Data-driven, yes. Proactive, not yet.

Programs such as EASA's Data4Safety[12] and DATAPP[13] initiatives highlight the potential of integrated, multi-source analysis to support this transition, but they also underscore how far current industry practice remains from fully predictive, intelligence-driven safety management[14].

[9]International Civil Aviation Organization (ICAO), Safety Management Manual (SMM), 4th ed., Doc 9859 (Montreal: ICAO, 2018).

[10]European Union Aviation Safety Agency (EASA), European Plan for Aviation Safety (EPAS) 2024–2026 (Cologne: EASA, 2023).

[11]International Civil Aviation Organization (ICAO), Evidence-Based Training Manual, Doc 9995 (Montreal: ICAO, 2013).

[12]European Union Aviation Safety Agency (EASA), "EASA Data4Safety", accessed 2026.

[13]European Union Aviation Safety Agency (EASA), "EASA DATAPP", accessed 2026.

[14]European Union Aviation Safety Agency (EASA), "Data4Safety: A New Approach to European Aviation Safety," accessed 2025.

The cost of missed opportunities

Missed opportunities are costly and the value locked in improved safety intelligence is both measurable and largely uncaptured.

Aviation's safety record is a real success, but it has come with an underappreciated economic downside: the industry has optimized for incident reduction without fully accounting for the value lost through operational variability, hidden disruptions and systemic blind spots that precede incidents. This reflects the historical limits of reactive safety systems that measure the cost of incidents after they occur, but not the cost of the chain of conditions that builds up and leads to them.

The direct economic impact of safety-related events is substantial and well-documented. Runway excursions are among the highest-consequence event categories tracked by ICAO and Authorities, which carry average direct costs exceeding $10 million per occurrence when hull damage, aircraft on ground (AOG) periods, investigation, and litigation are considered. Unplanned go-arounds driven by unstable approach conditions could generate direct operational costs in the range of €5,000–15,000 per event in fuel burn, crew time, slot reassignment, and connection disruption. Turnback events, unscheduled returns to departure airport, could cost between €20,000 and €80,000 per occurrence depending on aircraft type and stage length, excluding downstream schedule propagation.

Aggregated across the global commercial fleet, IATA estimates that accidents and serious incidents with material damage cost the industry in excess of $6 billion annually in direct and indirect costs. These figures, significant as they are, represent only the visible surface of a much larger economic problem.

The deeper and more consequential cost is structural, and it does not lie in the safety actions themselves, which are necessary, but in the upstream conditions that make them more frequent that they need to be. As an example, a go-around is not an inefficiency, it is the right response to an unstable approach.

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

The inefficiency is the chain of conditions (suboptimal planning, inadequate situational awareness, undetected environmental factors…) that made the approach unstabilised in the first place. Better safety intelligence does not eliminate the go-around as an effective barrier but might reduce its frequency by addressing root conditions earlier in the operational cycle.

These upstream conditions accumulate silently across fuel consumption, maintenance cycles, crew utilization, among others. These rarely appear in safety reports because they fall below reporting thresholds or are absorbed as normal operational variance. Fuel overconsumption linked to suboptimal approach profiles, reactive routing decisions, and unplanned holding patterns accounts for an estimated 3–7% of unplanned fuel expenditure in complex operational environments.

Maintenance cost acceleration driven by hard landings, off-envelope operations, and additional pressurization cycles appear in MRO ledgers months after the causal event, entirely disconnected from the safety record that generated it. For passengers, a single hub turnback can generate 50–200 misconnections in cascade: at a conservative average reaccommodation cost of €250 per affected passenger, a single event can produce €12,000–50,000 in direct passenger cost before reputational effects are considered.

Perhaps the most significant and least addressed cost is the systematic mispricing of risk in the insurance market. Current premium structures are calibrated on historical loss data and broad operator categories, with limited visibility into the actual operational risk profile of individual carriers. The result is a persistent cross-subsidy: operators with demonstrably stronger safety performance (e.g. better approach stability rates, lower precursor frequency, more mature safety management maturity) pay premiums that might not reflect their real risk position, effectively subsidizing operators with more risk exposure within the same market segment. Industry benchmarks suggest that the gap between current premiums and actuarially fair premiums for top-quartile operators ranges from 8 to 15% of annual insurance costs. For a regional operator with a fleet of 10–15 aircraft, this represents €200,000–600,000 per year in unrealized value, a direct consequence of information asymmetry between operators and the insurance market.

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

WHEN IT COMES TO SAFETY COSTS, THERE IS MORE THAN MEETS THE EYE

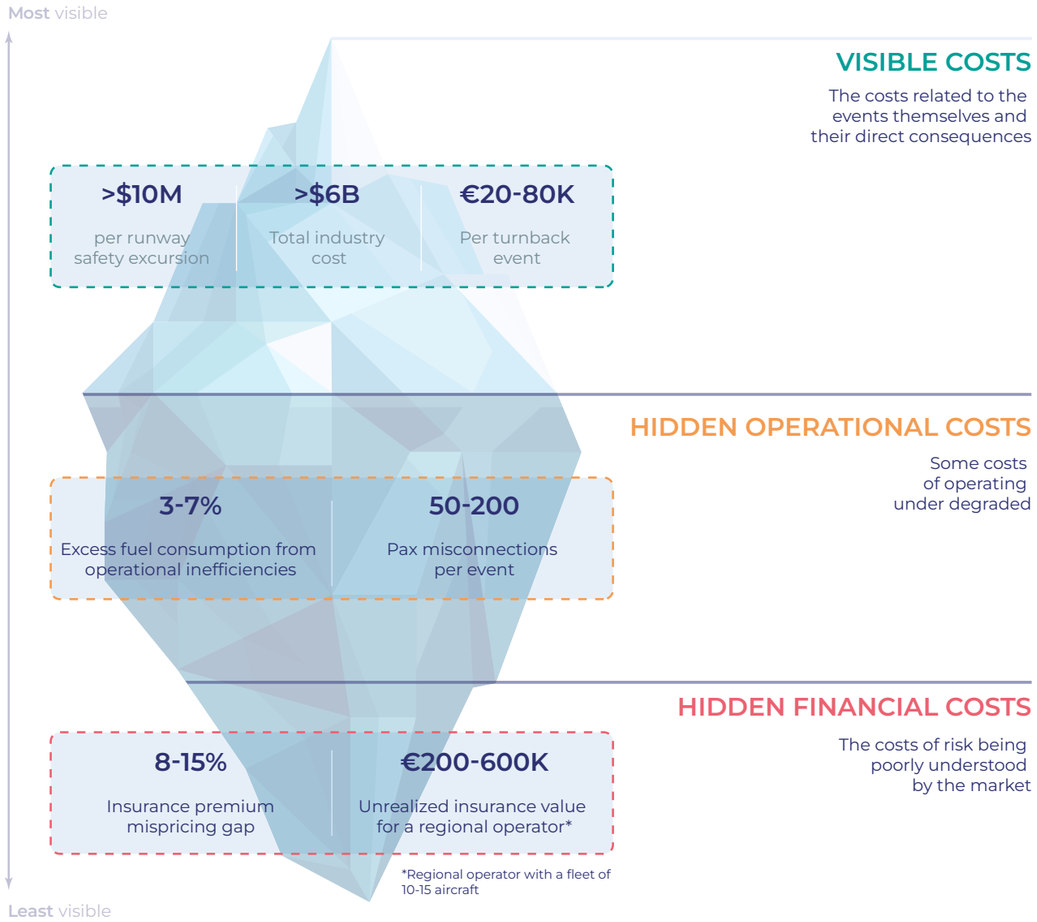

The value locked in improved safety intelligence is therefore not simply the avoided cost of incidents that do not occur. It is the aggregate of operational consistency recovered by addressing safety-relevant conditions earlier, before they escalate into events that demand a response. Earlier detection and systemic understanding do not eliminate the need for sound operational judgement but they might reduce the frequency with which that judgement is tested under pressure. The question for operators, airports, and their insurance and financial partners is no longer whether this value exists, but whether their current analytical capabilities are sufficient to claim it.

SAFETY INTELLIGENCE: A NEW PARADIGM IN AVIATION SAFETY

Unlocking actionable insights from data is the way forward for aviation safety

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

Defining Safety Intelligence

Safety Intelligence is formally defined by ICAO as "the outcome of a systematic and structured process of collecting, analyzing and interpreting safety data and safety information, with the objective of supporting data-driven decision making"[15]. This definition establishes Safety Intelligence as the result of transforming raw safety data into meaningful insights that can effectively inform safety-related decisions, rather than as the data or analytical processes themselves.

Beyond the formal definition, Safety Intelligence is better understood as an organizational capability than a technical function. Advanced tools and platforms matter but they are enablers, not the capability itself. Safety Intelligence cannot be reduced to an IT system that can be purchased or deployed in isolation. Instead, it emerges from the combination of data availability, governance frameworks, analytical capabilities, organizational processes and a culture of collaboration and learning across stakeholders. Its value lies not only in analyzing safety data, but in integrating information from multiple actors and domains, identifying patterns and precursors, and translating those insights into informed actions that support operational, tactical and strategic decision making.

Figure 1 shows the hierarchy of data, information, and intelligence, illustrating Safety Intelligence as a multi-source interpretation that allows for understanding patterns and supporting decision making.

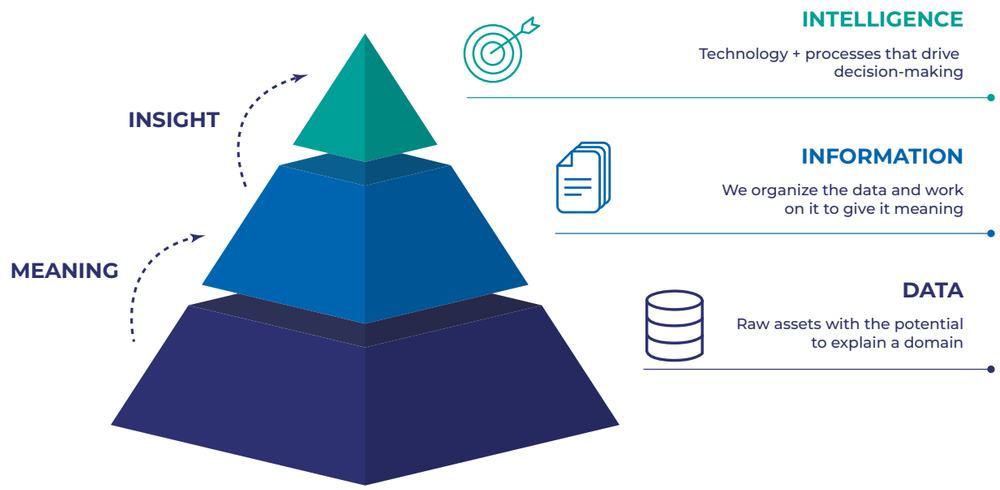

Figure 1 – Safety Intelligence emerges from increasing levels of abstraction and analysis of relevant data

[15]International Civil Aviation Organization (ICAO), Safety Intelligence Manual, Doc 10159 (Montreal: ICAO, 2025).

Safety Intelligence as an enabler of safety management processes

By its very nature, Safety Intelligence is a central enabler of modern SMS and State Safety Programmes (SSPs). By transforming integrated safety data into actionable, evidence-based insights, it supports more effective hazard identification, risk profiling, and the definition of meaningful safety objectives and performance indicators. In doing so, it underpins risk based oversight and other monitoring strategies, strengthening the link between safety data, risk assessment and decision making, and improving safety performance in increasingly complex operational environments.

It operates within a continuous cycle: analyses generate insights that inform safety decisions and action plans, while the outcomes of those actions produce new data that feeds back into the system, gradually enhancing the organization's ability to anticipate emerging risks. As this cycle matures, organizations build more robust Safety Intelligence capabilities that reinforce both safety performance and proactive risk management.

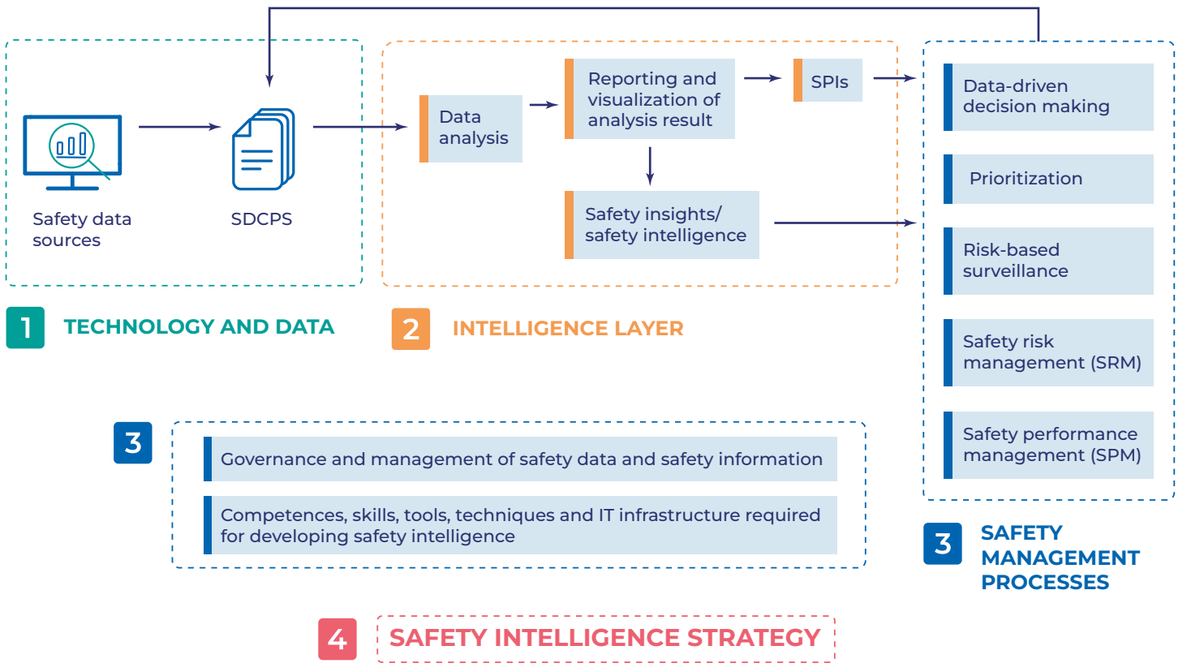

Figure 2 – Safety Intelligence as an enabler of safety management processes

ALG

Figure 2 describes this feedback loop in detail, showing how four interdependent layers form the backbone of a functioning Safety Intelligence ecosystem:

1

Technology and data

At the core of safety intelligence lie the data collection processes and underlying technology that drives insight generation.

2

Intelligence layer

The existing data is then analzed to generate relevant safety information and insights/intelligence.

3

Safety management processes

The results of the data analysis inform the processes that form the backbone of SMS, including data governance, data-driven decision making, and the management of the relevant skills and infrastructure necessary to support the safety intelligence ecosystem.

4

Safety intelligence strategy

Overarching the entire ecosystem is the safety strategy, which contains a roadmap to transform the system from its current state to the target vision.

Developing that ecosystem requires a structured approach that clarifies how decisions are made, which safety questions matter most, and what data and analytical capabilities are needed to answer them. There is no single organizational model suitable for all contexts: responsibilities may be centralized, distributed or supported by external partners, but effective Safety Intelligence functions should remain sufficiently independent to limit confirmation bias while being collaborative enough to integrate insights across technical and organizational boundaries.

Safety Intelligence in practice

In practice, Safety Intelligence capabilities can be observed at different levels of maturity, which can be broadly broken down into three levels, as shown in Figure 3:

Traditional reactive approaches, focused on incident investigation and compliance reporting, provide limited value in complex operational environments.

Proactive approaches are emerging, enabling earlier detection of safety risks through integrated data analysis.

Fully predictive and prescriptive capabilities remain largely aspirational, but represent a realistic target as governance, technology and organizational readiness evolve.

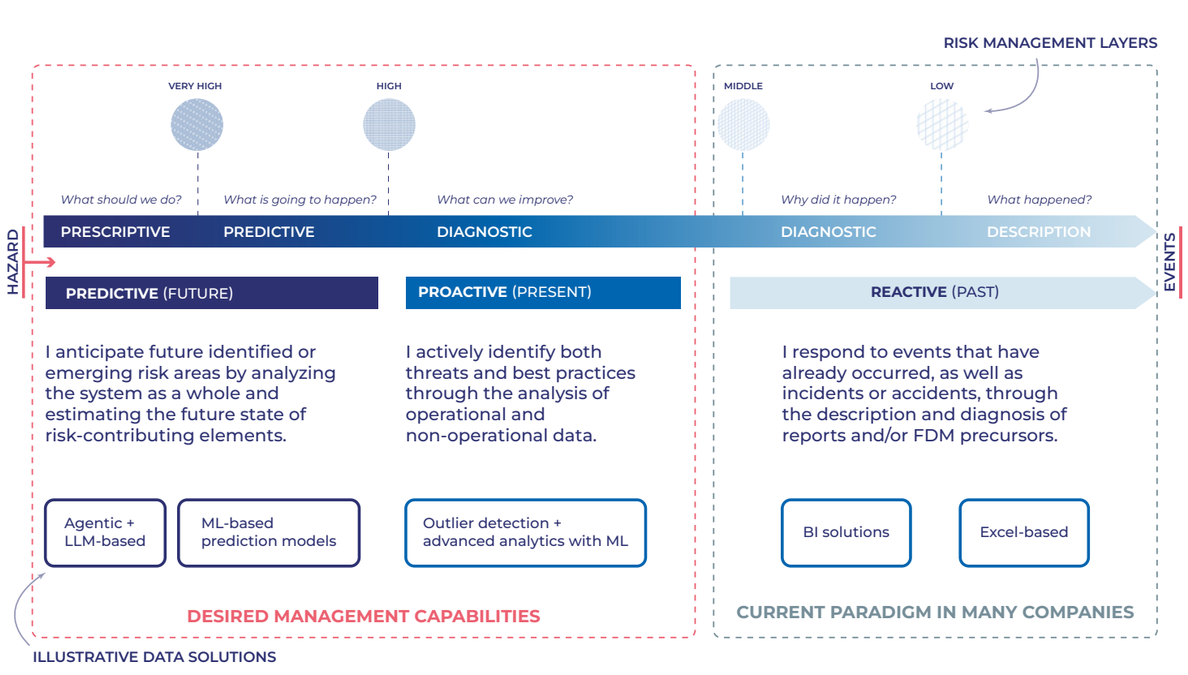

Figure 3 – Safety Intelligence emerges through the transition from reactive to predictive approaches

To understand how these different levels of maturity can be used in the context of a specific safety incident, consider the example of a long landing leading to an inability to stop within the expected landing distance. One can characterize such as incident using the so-called bow-tie model, shown in Figure 4.

The incident sits at the center of the bow-tie. On the left are the precursors that increase the likelihood of the event: incorrect deployment of stopping devices, unstable energy management on approach, environmental conditions. Between precursors and the event itself are barriers — standard operating procedures, go-around decision criteria, crew monitoring. On the right are the consequences: runway excursion, collision with obstacles beyond the runway end. The bow-tie is a powerful lens for comparing how different levels of analytical maturity affect the ability to prevent, detect, and mitigate such a scenario.

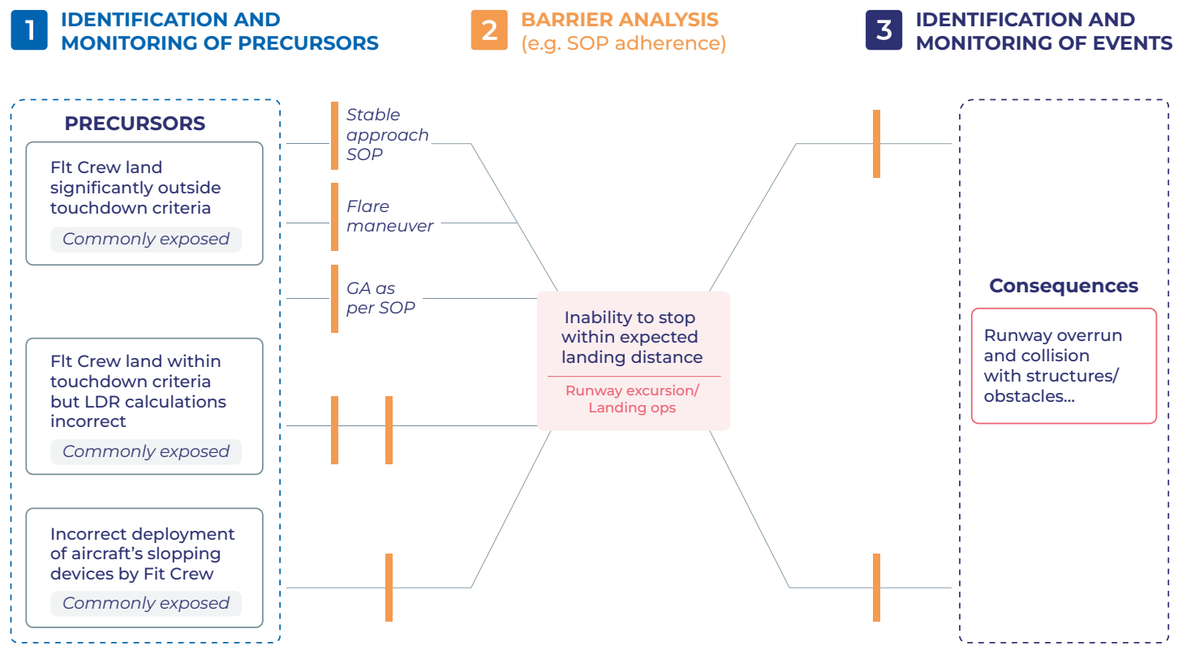

Figure 4 – Inability to stop within the expected landing distance through the lens of a bow-tie model

Level 1 – Reactive approach

Reactive Safety Intelligence focuses on understanding events that have already occurred, consolidating evidence from reports and operational data and supporting descriptive and diagnostic analysis to inform corrective actions. At this maturity level, value depends heavily on consistent data capture, integration and efficient reporting cycles. Currently, many organizations remain this paradigm, often relying on Excel-based or basic BI solutions to track these outcomes.

Using the bow-tie model of Figure 4, reactive analysis starts from the event outcome and reconstructs the pathway: which precursors were present, how barriers performed, and how the sequence progressed towards the top event ("inability to stop within expected landing distance") and its consequences (runway excursion). The objective is to identify what happened and why, and to translate findings into targeted mitigations.

Analytical capabilities at this level: automated identification of events and precursors from FDM and reports; root-cause diagnostics and barrier-focused analysis; SPI dashboards and periodic safety performance reporting; standardized workflows for investigation outcomes and action tracking.

Level 2 – Proactive approach

Proactive Safety Intelligence moves beyond isolated occurrences and focuses on the continuous assessment of normal operations. Its objective is to identify degradation signals, patterns and improvement opportunities before they materialize into incidents. This shift requires broader data integration, repeatable analytics, and governance that enables cross-domain interpretation.

In bow-tie terms, proactive Safety Intelligence shifts attention to the left side of the pathway: using diagnostic analytics to monitor threats and contributing factors that increase exposure to long-landing scenarios. This shift is shown in Figure 5. Instead of analysing only the cases where the main event took place, proactive analytics track how often precursors appear, how barrier health evolves over time, and which combinations of factors systematically increase risk. This enables earlier prioritization of preventive actions across training, procedures and operational controls.

1 MONITORING OF OPERATIONAL INDICATORS + *DYNAMIC DEGRADATION ALERTS*

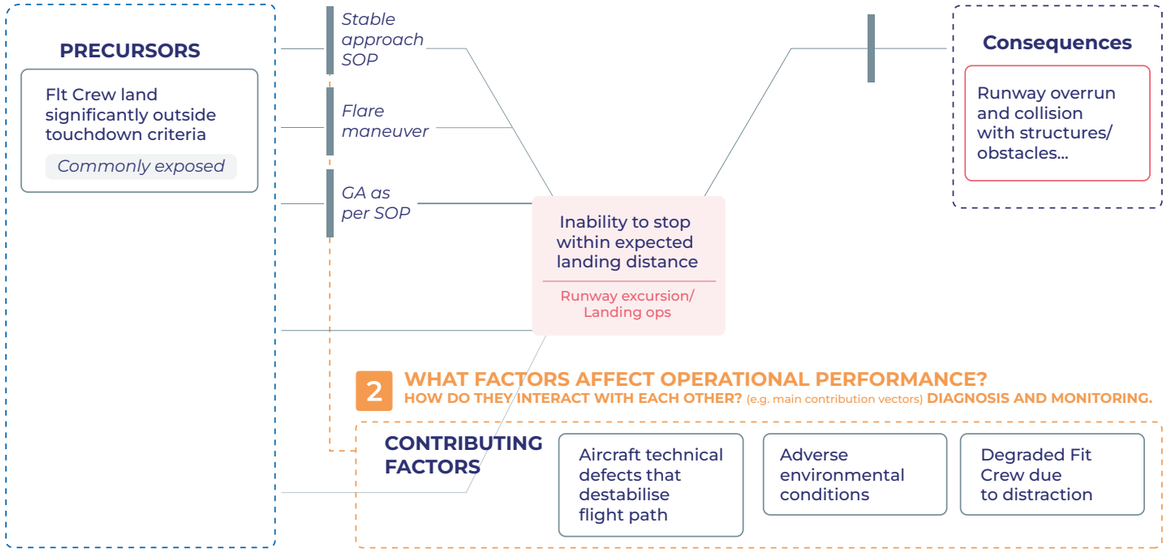

Figure 5 - Proactive approaches extend the basic bow-tie model with the inclusion of contributing factors that affect existing barriers between the precursors and the safety incident itself

Analytical capabilities at this level: continuous monitoring of leading indicators and safety margin degradation; cross-source correlation to identify contributing-factor patterns; outlier detection and trend analytics at fleet, airport, and system level; effectiveness monitoring of mitigations.

ALG
Safety Intelligence in aviation:
The new reality for prescriptive safety management

Level 3 – Predictive approach

Predictive Safety Intelligence aims to anticipate emerging risk areas by forecasting how contributing factors and barrier performance may evolve over time. While less widely deployed today, it is a realistic target enabled by mature data platforms, machine learning and robust governance. Its value lies in translating forecasts into timely preventive decisions, not in prediction for its own sake.

Applied to the same bow-tie, predictive approaches forecast the evolution of contributing factors, predicted weather, traffic constraints, technical reliability signals, and the likely degradation of key barriers, such as training effectiveness indicators and SOP adherence trends. Figure 6 illustrates this. The output is forward-looking risk profiles that trigger alerts when projected exposure crosses defined thresholds, enabling preventive intervention before safety margins erode.

Analytical capabilities at this level: forecasting models for contributing factors; probabilistic models for top-event risk and scenario-based risk profiles; dynamic alerting and thresholding linked to preventive actions; decision support “copilot” concepts with human-in-the-loop control.

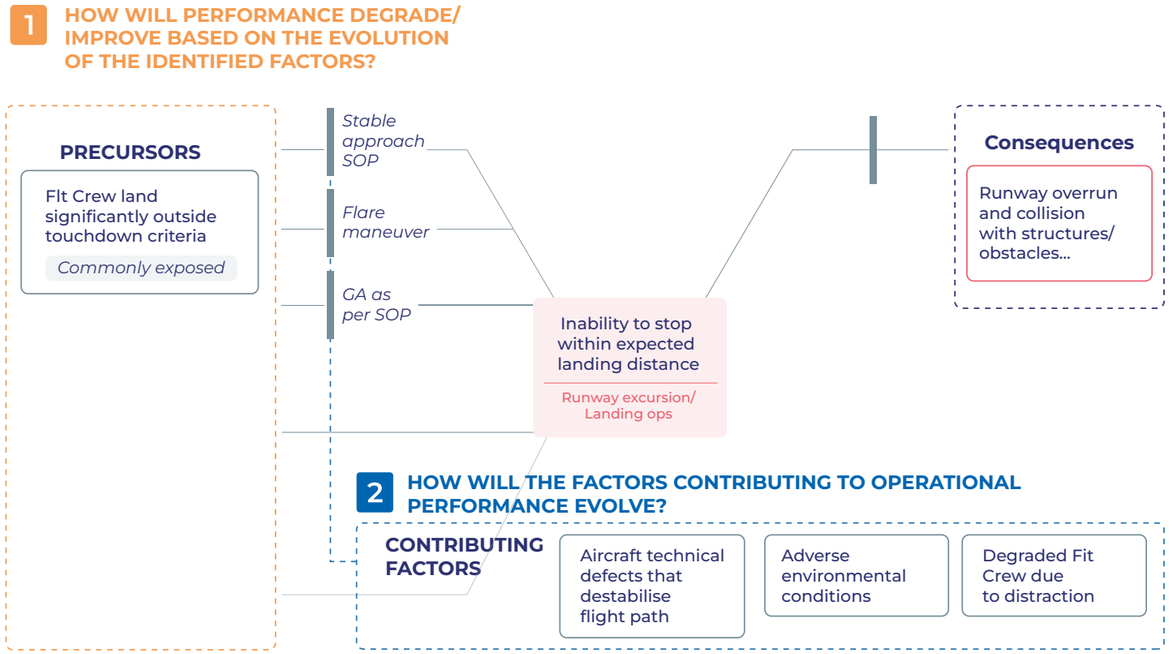

Figure 6 - Predictive approaches monitor the evolution of contributing factors and quantify the likely degradation of the relevant barriers, yielding an a-priori estimate of the probability of the safety incident taking place

Reference use cases

Safety Intelligence dashboarding and digitalized safety reviews

A representative example of reactive Safety Intelligence is the deployment of integrated dashboard environments that support recurrent safety reviews and internal safety action groups. These environments consolidate Safety Performance Indicators (SPIs) across priority risk areas and allow structured drill-down into individual flights or events for root cause analysis. Within collaborative initiatives such as EASA's Data4Safety programme, this capability is enabled by pooling safety relevant data (FDM, occurrence reports, operational data, meteorology and surveillance) into a shared analytical infrastructure that produces metrics dashboards, blind benchmarking views and directed studies.

Building on these principles, ALG has implemented such dashboarding and digital safety review environments for airlines, ANSPs and airports, helping them harmonize and digitalize safety assessments and shorten the feedback loop between occurrence analysis, corrective actions and performance monitoring[16].

[16]EASA – Data4Safety – 2024. An exciting year ahead for Data4Safety.

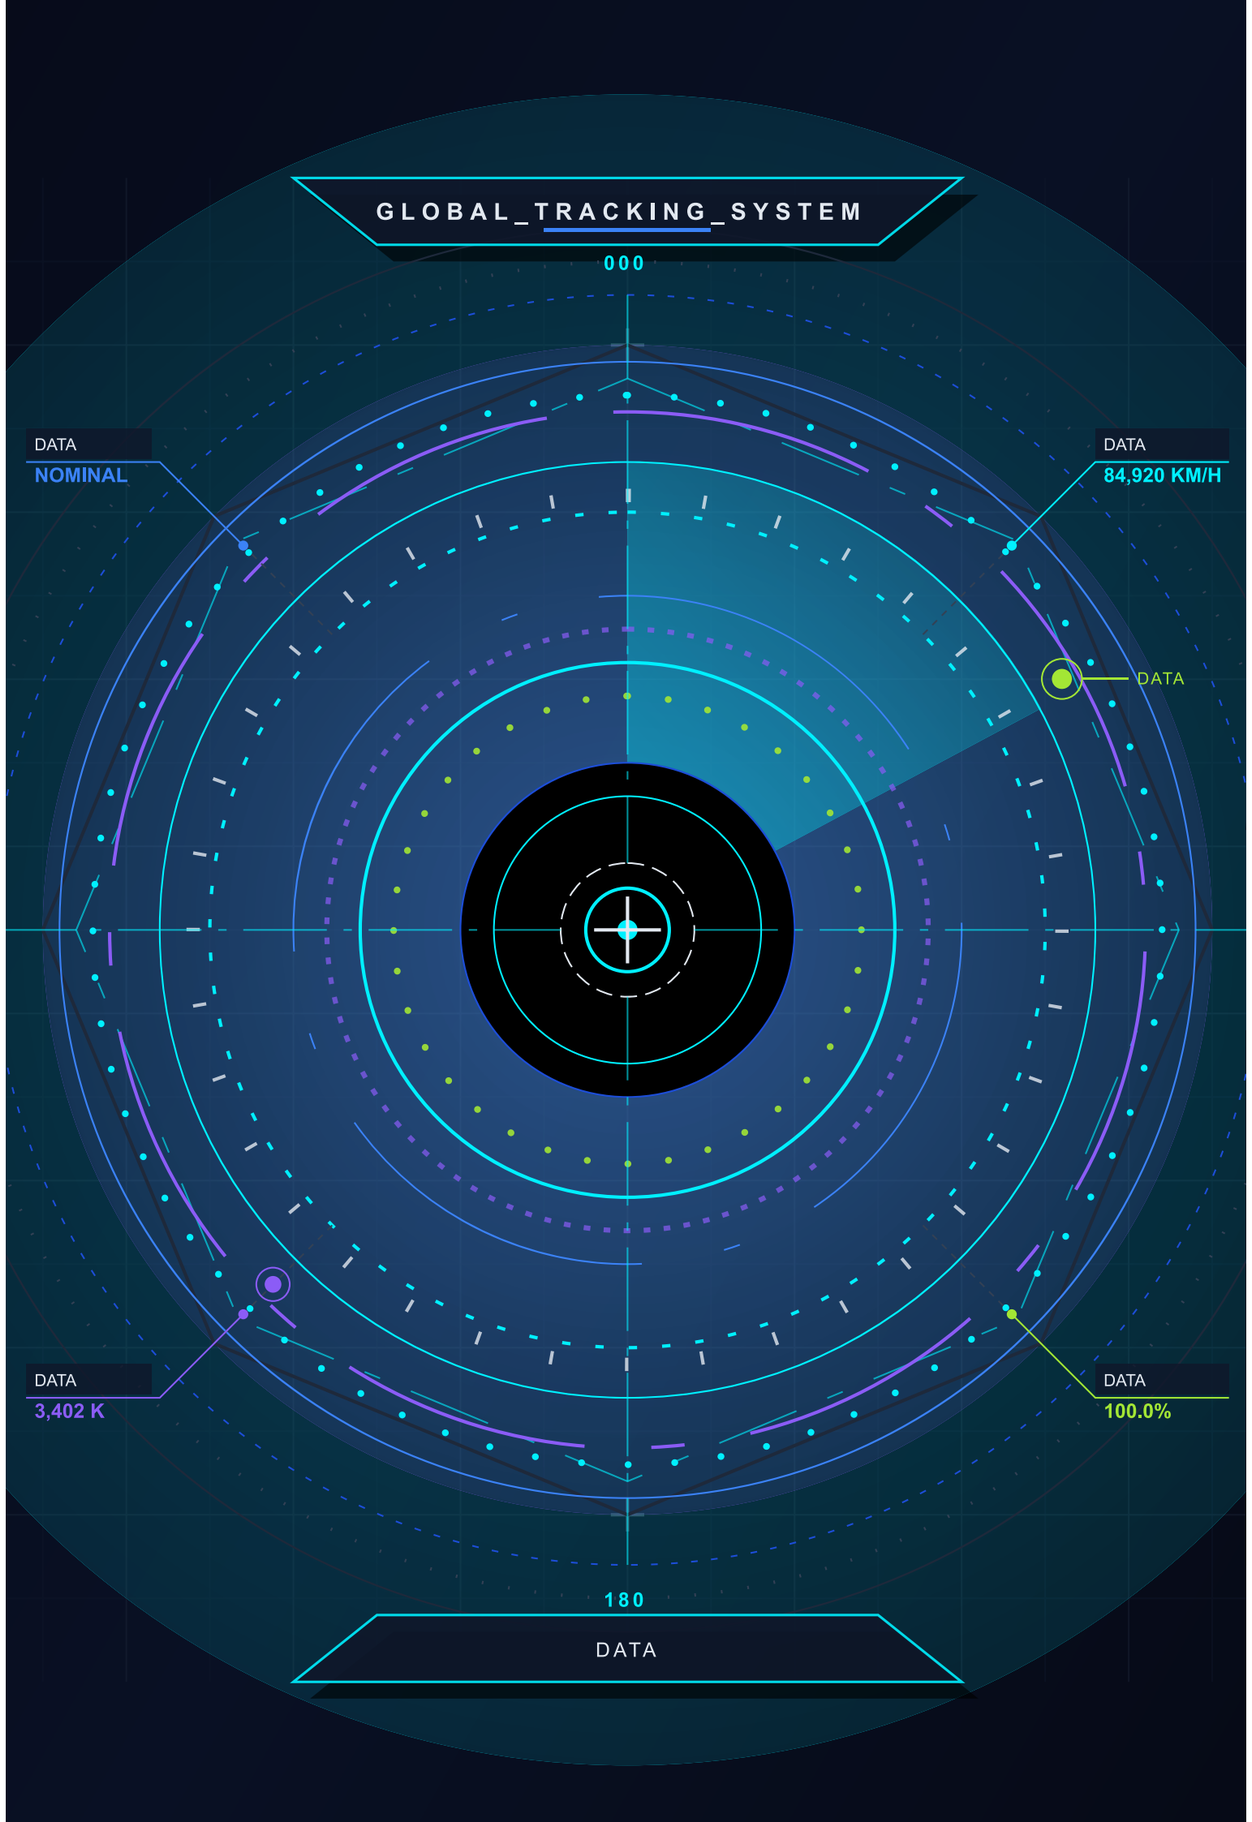

Proactive Safety Intelligence – Approach path management at scale

The development and scaling of proactive Safety Intelligence capabilities is particularly visible in collaborative analytical frameworks focused on normal operations. A prominent example is approach path management, as explored in Data4Safety directed studies on unstable approaches and go-arounds. In this use case, FDM and surveillance data (e.g. ADS B) from multiple operators and airports are analyzed using harmonized detection algorithms and criteria to understand the drivers and consequences of unstable approaches at the system level.

ALG has designed and deployed such approach path management analytics, enabling clients to move from isolated event analysis to systematic identification of structural patterns in vertical and lateral profiles, energy management, environmental conditions and operational constraints. The resulting insights support targeted improvements in training, procedures and airspace or runway use, strengthening safety margins across fleets and airports while providing a concrete, operational example of proactive Safety Intelligence in practice.

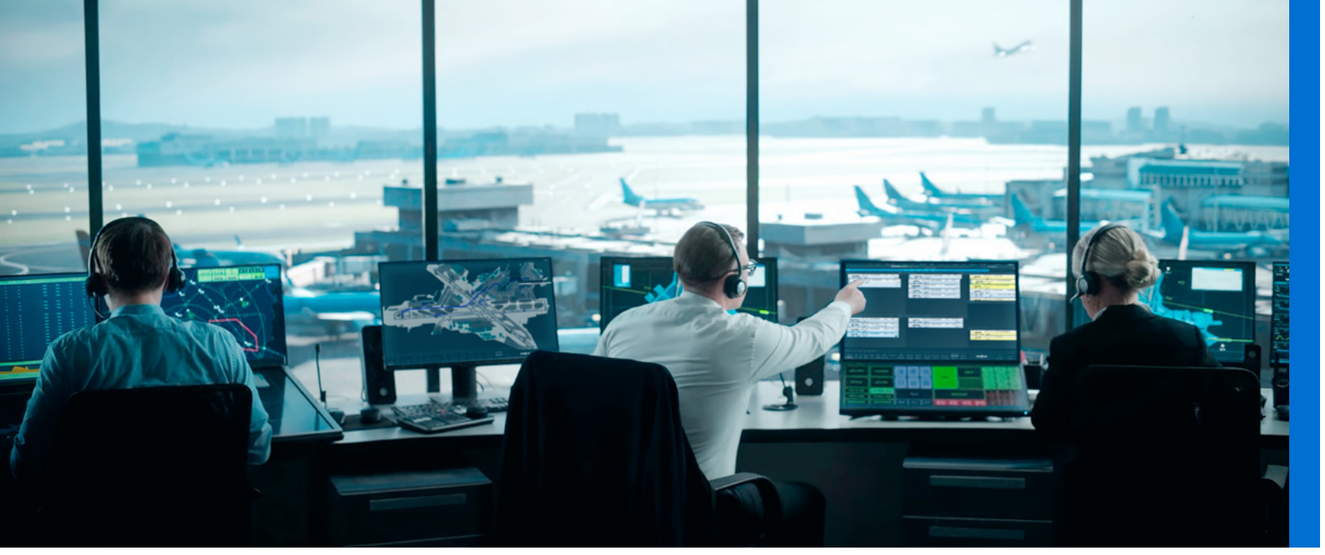

THE PATH AHEAD: CHALLENGES, BARRIERS, AND THE WAY FORWARD

How aviation can evolve from reactive reporting to intelligence-driven decision making

Predictive Safety Intelligence – AI-enabled safety briefing and decision support

Predictive Safety Intelligence represents the next step, anticipating emerging risk by estimating how contributing factors and barrier performance may evolve. Industry trends point clearly in this direction, and the concept is technically feasible today.

A forward-looking example of what predictive safety intelligence could look like in practice are AI-enabled safety briefing and decision-support tools for safety analysts. Such tools would continuously synthesize safety-relevant signals and translate complex analytical outputs into tailored insights. For example, such an agent would be able to identify, a-priori, specific flights that are at risk of a specific safety incident taking place based on its understanding of contributing factors and their consequences. This information would then be relayed to the relevant stakeholders as needed (e.g. airline/airport safety managers, flight crew, air-traffic controllers, airport operators, etc.), flagging risky operations and increasing safety oversight before an incident takes place.

Rather than replacing expert judgement, these copilot capabilities provide forward-looking risk profiles, scenario-based what if analyses and prioritized recommendations, helping organizations focus limited expert capacity on the most critical issues while maintaining human-in-the-loop control over safety decisions.

Why Safety Intelligence is hard

The case for Safety Intelligence is clear. The path toward it is not. The barriers are real and worth naming honestly, because understating them is how implementation projects fail.

Technical complexity

From an analytics perspective, aviation's strong safety record creates a real modeling challenge: accidents are rare, and serious outcomes are sparse. Pure "accident prediction" is often the wrong target. The practical path is to model precursors, abnormal patterns, and barrier degradation using large-scale operational data and then translate those signals into risk-informed action. Traditional Machine Learning solutions, coupled with Agentic solutions might come at hand. However, this adds an additional technical complexity: any AI-supported insight must be explainable enough to earn trust in operational contexts. If a model behaves like a black box, it risks being ignored by professionals trained in evidence-based decision making and accountable for outcomes. This is why aviation AI roadmaps emphasize human-centric, trustworthy approaches and why anomaly/precursor detection is often a more tractable starting point than outcome prediction[17].

[17]Anomaly Detection in Aviation Data using Extreme Learning Machines. Vijay Manikandan Janakiraman. David Nielsen. 2016

ALG

Safety Intelligence in aviation:
The new reality for prescriptive safety management

The governance and legal paradox

At the core of this challenge lies an inherent tension between transparency and exposure. High-quality safety insight depends on the willingness of stakeholders to share granular operational data and candid occurrence narratives under a "Just Culture" mindset. However, legal frameworks and investigative practices differ across jurisdictions, and the boundary between safety confidentiality and judicial discovery is not always clear. This creates a chilling effect: the more detailed the reporting, the higher the perceived legal risk. As a result, organizations may restrict access, de-identify excessively, or avoid sharing altogether, limiting system-wide learning and reducing the value of collaborative safety intelligence[18,19].

Interoperability and change management are the "silent blockers"

Even when data can be shared legally and models can be built technically, many Safety Intelligence initiatives fail in practice due to friction: incompatible taxonomies, inconsistent data quality, missing context, and unclear ownership. On the human side, changing routines is hard: who reviews the insight, who decides, how quickly actions are taken, how accountability is handled, and how learning is fed back into existing processes are valid questions to raise. Without clear decision loops and governance, Safety Intelligence becomes "more dashboards" rather than measurable safety improvement[20].

[18]Just Culture. Eurocontrol.
[19]Just Culture Manifesto. Skybry.
[20]Artificial Intelligence Roadmap 2.0. EASA. May 2023.

What needs to change

In practical terms, three changes matter most, being more organizational than technical:

- **Moving from local confidentiality to governed collaboration**

Safety Intelligence scales when stakeholders can collaborate without eroding “Just Culture” principles. This requires explicit governance: data-sharing agreements, role-based access, de-identification standards, and mechanisms to protect sensitive narratives while preserving analytical value. The goal is not “open data” but trusted, governed sharing that increases learning without increasing fear of blame.

- **Building interoperable safety information to overcome fragmented datasets**

Getting the full risk picture depends on cross-domain insights, which are entirely dependent on using a common language. Without a shared taxonomy and minimum metadata standards, the burden of manual reconciliation becomes prohibitive. Industry initiatives show the direction of travel: integrating multiple data sources into structured environments that support benchmarking and trend identification. These efforts reinforce the need for standardization as a foundation for scalable intelligence[21].
Recent advancements in AI will ease the standardization work from both historic and future safety data.

- **Shifting from periodic reporting to risk-based decision-making and oversight**

This one is the most operationally immediate. To turn insights into impact, organizations need decision protocols: thresholds, alerting logic, escalation paths, and clear ownership of actions. At the ecosystem level, oversight models are also evolving toward approaches that are more risk and performance informed, which aligns naturally with Safety Intelligence maturity. Rather than framing this as “regulatory relief”, the constructive framing is risk-based oversight alignment: organizations that demonstrate strong internal intelligence and transparent safety performance enable more targeted, effective oversight focused on risk areas.

Finally, Safety Intelligence should be framed as part of a broader systemic resilience agenda. Modern aviation risk is interdependent: safety, information security, human performance, and socio-economic pressures interact in ways that no single domain can manage alone. The 2026 European Plan for Aviation Safety (EPAS) explicitly highlights these interdependencies. Safety Intelligence, done well, is the integrating layer that makes systemic resilience operational[22].

[21] IATA. Global Aviation Data Management (GADM).

[22] EPAS. European Plan for Aviation Safety. Volumen II. 2026 Edition.

ALG

Safety Intelligence in aviation:
The new reality for prescriptive safety management

A practical, realistic roadmap

The transition toward Safety Intelligence is evolutionary rather than disruptive. A realistic roadmap can be articulated across three horizons:

HORIZON 1
(1–2 YEARS)

DATA MATURITY AND RELIABLE REACTIVE INTELLIGENCE

Focus on making reactive intelligence efficient and trustworthy: improve reporting timeliness, consolidate key data sources, and operationalize a “single source of truth” for SPIs and investigation workflows. Typical outcomes include faster diagnosis of events and more consistent barrier-focused learning across the organization. This horizon is about foundations.

HORIZON 2
(2-3 YEARS)

PROACTIVE, CROSS-DOMAIN INTELLIGENCE AT SCALE

Expand from event-centric analysis to continuous monitoring of normal operations: leading indicators, precursor patterns, barrier effectiveness tracking, and cross-source correlation. Where full data pooling is constrained, progress can still be achieved through architectures that enable sharing of anonymized insights and aggregated risk patterns across stakeholders (airlines–airports–ANSPs), strengthening systemic learning without forcing full data centralization.

HORIZON 3
(3-5+ YEARS)

PREDICTIVE DECISION SUPPORT AND SYSTEMIC RESILIENCE

As governance, interoperability and analytical maturity increase, Safety Intelligence can evolve toward forecasting and prescriptive support: forward-looking risk profiles, dynamic alerting linked to preventive actions, and more integrated “system-of-systems” views (often described through digital-twin-like concepts). This should be positioned as a long-term vision: achievable progressively, but dependent on strong foundations and human-in-the-loop decisioning.

CONCLUSIONS

The aviation industry has spent decades perfecting its ability to learn from accidents. That investment has paid off in ways that are measurable and genuinely significant. But the next wave of safety improvement will not come from better accident investigation. It will come from the ability to read the signals that precede accidents and act on them before they become events.

That is what Safety Intelligence makes possible. Not as an abstract concept, and not as a technology purchase. As a structured organizational capability that connects data across domains, generates insight that would otherwise stay hidden, and puts that insight in front of decision-makers who can use it in time to matter.

The building blocks are available. The data exists. The analytical methods are mature enough to deploy. What has been missing, in most organizations, is the governance, the integration, and the organizational will to move from compliance-driven reporting toward intelligence-driven decision making. Those are solvable problems, way harder than technical ones in some ways, but entirely within reach for aviation stakeholders who choose to prioritize them.

The organizations that move early will not just have safer operations. They will have more efficient ones. They will be better positioned in insurance markets. They will have a demonstrable, data-backed answer to the question that regulators, partners, and investors increasingly ask: not “have you had any incidents?” but “do you understand your risk?”.

Advancing in this direction requires the integration of safety data, analytical capability, and operational decision-making into a coherent and sustained capability. While the data and methods are increasingly available, the key differentiator will be organizational alignment, governance, and the ability to act on insight in a timely and consistent way.

ALG

Núria Alsina Pujol
MSc in Aeronautic Engineering, Principal
nalsina@alg-global.com

Antonio Cabeza Doña
MSc in Aeronautical Engineering, Senior Engagement Manager
acabezad@alg-global.com

Pablo González Martinez
MSc in Aerospace Engineering, Consultant
pgonzalezmar@alg-global.com

alg-global.com