

THINK PAPER

MARCH 2026



AIR NAVIGATION SAFETY PERFORMANCE: A DRIVE TO MOVE BEYOND COMPLIANCE METRICS

SHAPING
OUR
FUTURE
SKIES

canso.org

Abstract

This CANSO Think Paper explores the limitations of traditional compliance-based safety metrics and examines the potential of Time to Conflict (TTC) as a complementary, risk-based approach to safety management in air navigation.

Historic compliance rules and procedures have been very useful in improving aviation safety worldwide; however, recent advancements in computing capability have resulted in the ability to determine the true risk of encounters between aircraft, even when aircraft meet compliance rules. Adoption of Time to Conflict (TTC) algorithms and processes will permit Air Navigation Service Providers (ANSPs) to further improve safety through the identification of latent safety risk events and the mitigation of risk in such events.

With the reduction of mid-air collisions resulting from improved aviation safety globally, losses of separation have traditionally been used as a proxy for mid-air collision risk. ANSPs record and monitor the number of instances where losses of separation occur and attempt to quantify the severity of each event. These are referred to in this paper as “compliance metrics”, describing how well required separation is maintained to demonstrate safety performance. Compliance rules have been used since the inception of Air Traffic Control procedures and have served the flying public well as a proxy for risk. They are used to indicate that a given rule or standard has been met and that the operation is therefore considered safe.


Compliance parameters are increasingly losing their ability to assure that risk is being managed to an acceptable level, as defined by the International Civil Aviation Organization (ICAO). To provide such assurance, ANSPs need a better understanding of the risks involved in air navigation. This paper discusses some of the challenges associated with introducing a more suitable metric based on the TTC algorithm.

With improved computing capabilities, TTC algorithms and procedures can now be developed and implemented to provide valuable insight into latent safety risk. This includes both compliant events with relatively high levels of risk and non-compliant events with comparatively low levels of risk. ANSPs may therefore benefit from adopting TTC algorithms and processes that augment existing compliance-based rules.

The adoption of TTC processes and algorithms also assists organisations by prioritising events for more detailed analysis. TTC provides additional information on aircraft encounters that can be used to better understand operational risk and to introduce additional barriers to prevent similar events from reoccurring. This can be particularly beneficial for ANSPs facing resource constraints, enabling a more focused and efficient approach to safety analysis and review.

Finally, this Think Paper also serves as a call for collaboration to further develop risk-based safety management concepts in air navigation.

Index Terms: Air navigation, separation standards, time to conflict, risk-based safety management.



Introduction

The role of an Air Navigation Service Provider (ANSP) is to regulate efficiency and ensure safety. Aviation professionals are guided by the principle of “safe, orderly, and expeditious” operations, commonly emphasised in the aviation industry in that specific sequence. It can be argued that the primary safety concerns for ANSPs relate to in-flight and ground collisions, wake vortex accidents, and collisions with obstacles on the ground. These events are highly significant and attract global visibility, as they are often accompanied by breaking news coverage.

Fortunately, such events are extremely rare, and ANSPs therefore rely on other events as proxy or secondary parameters. These include losses of separation, runway incursions, altitude deviations, airspace infringements, route diversions, interception above glide path, loss of communication, and similar occurrences. The aviation industry also applies additional means to assure the public, ANSPs, safety regulators, and safety professionals that safety is being maintained, using metrics such as:

- Audits;
- Maturity scores of Safety Management Systems (SMS);
- Safety-II;
- Surveys, and;
- Normal Operations Safety Surveys.

Most of these secondary parameters measure outcomes against a defined standard (e.g. ICAO). For losses of separation, this may be less than five or three nautical miles (NM) and 1,000 feet. For runway incursions, it may be defined as an aircraft or vehicle entering a pre-defined Runway Protected Area (e.g. crossing a hold-short line). When a standard is breached, a safety event is recorded. These can be classified as compliance metrics — indicators of whether a given standard has been met or whether non-nominal performance has occurred.

The Problem

The industry is entering a phase in which these metrics are becoming insufficient, offering diminishing insight as non-compliant events become less frequent. A further complication is that such events primarily describe past outcomes. While lessons from past events are essential, the causal factors leading to non-compliance have already occurred, whereas risk exists in the future. A metric showing no accidents over several years is reassuring, but it provides limited insight into the risks managed on a daily basis.

It is widely accepted that some compliant encounters (i.e. no loss of separation) may still carry the potential for adverse outcomes, while some non-compliant events may involve relatively low risk. In high-risk maneuvering environments, such as Terminal Maneuvering Areas (TMAs), safety depends on the combined performance of air traffic controllers, pilot responses, and critical systems (e.g. planning tools, half-duplex radio, Short-Term Conflict Alert (STCA), and the Traffic Collision Avoidance System (TCAS)). Reliance on multiple interacting safety barriers in such environments has led to an industry consensus that measuring levels of non-compliance alone does not fully or accurately capture overall system resilience.

As a result, there is a growing need for safety performance metrics that provide greater insight into the actual risk of airborne encounters, rather than serving solely as measures of compliance.



The Case for TTC

Over the past decade, rapid advances in aviation concepts and capabilities have added further complexity to an already complex system. Globally, aviation is in a period of discovery and rapid innovation, incorporating new entrants, control concepts, and business models, including Unmanned Aircraft Systems (UAS), Advanced Air Mobility (AAM), electric vertical take-off and landing (eVTOL) aircraft, and commercial space operations. These developments challenge the status quo.

As of 2024, the Federal Aviation Administration (FAA) estimates an active fleet between 843,000 and 2,833,000 small UAS (weighing more than 0.55lbs and under 55lbs) with forecasted growth in 2029 reaching between 920,000 and 3,100,00 small UAS. For micro UAS (weighing under 0.55 lbs), the FAA estimates the fleet size to be 1.35 million, with a slight decline expected to 1.28 million by 2029. In contrast, substantial growth is anticipated in the large UAS fleet (weighing more than 55 lbs); the FAA estimates an active fleet of 4,314 large UAS in 2024, with this number rapidly increasing over tenfold to 44,740 by the end of 2029. Additionally, remote pilot certifications are expected to grow from 406,000 in 2024 to 497,000 in 2029. Using an unconstrained forecast, the FAA forecasts a fleet of four Advanced Air Mobility (AAM) vehicles performing over 40,000 flights during the first year of service, with these numbers expected to increase substantially to 283 AAM vehicles performing 2.8 million flights within five years.

Developments in these areas have laid the foundation for the continued evolution of the aviation domain. The resulting environment is expected to be mixed and transitional, with users exhibiting diverse capabilities and operational needs. This evolution, combined with the ANSP's responsibility to ensure safety in areas where rules and standards may be absent or immature, highlights the need for risk-based safety management. Likelihood and severity data will be essential to assess the feasibility and risk of integrated UAV operations, with collision risk being a critical component.

As this evolution continues, global engagement and alignment will be increasingly important, particularly with regions investing heavily in new entrants, including parts of Africa, Asia-Pacific (e.g. Singapore and South Korea), and the United States. Collaboration is needed to support the development or adaptation of SMS processes and the analysis of collision risk between crewed and uncrewed operations. Accordingly, it is

imperative that ANSPs develop and operationalise tangible methods to generate safety intelligence capable of addressing current and emerging risks in a fast-paced and dynamic aviation environment, while meeting public expectations for continued high levels of safety.

Regulatory Aspects

Risk-based safety metrics cannot replace compliance parameters, as rules remain essential safety tools that ensure the use of repeatable and well-established best practices. However, risk-based safety metrics can serve as complementary data points that provide improved insight into actual safety risks.

If the scoring of TTC encounters is combined with compliance metrics for each encounter (i.e. potential or validated losses of separation) to identify the relationship between the two, this can provide a useful basis for understanding compliance issues across other TTC encounters. The implementation of such combined scoring can also support the relative prioritisation of TTC encounters and enable comparison with similar encounters of varying scores that did not result in non-compliant or higher-risk event.

TTC algorithms and processes provide a means of relative prioritisation across TTC encounters. These algorithms estimate aircraft trajectories without considering intervention by pilots, controllers, or other safety barriers. As such, they can help identify isolated or systemic situations (e.g. airspace design) in which aircraft are guided into trajectories that may be considered higher risk and where safety relies on the effectiveness of existing barriers (e.g. procedural measures).

Analysis of TTC encounter contextual information, and their outcomes, may help clarify which encounters appear to be higher risk or non-compliant compared with similar events, and why. Quantifying aggregate risk and the resilience of safety barriers for repetitive or isolated encounters may also provide further insight into TTC encounters that are higher risk or non-compliant.

Insights from these processes may contribute to increased system safety resilience and a deeper understanding of how to avoid placing aircraft in such situations in the first place. Comparing projected risk with actual outcomes could support future refinement of compliance rules for specific operational contexts.

Risk-based safety metrics can also address operational areas for which no formal standards currently exist, and where compliance alone is therefore insufficient. If regulatory compliance is the sole focus, data on the complexity and risk characteristics of certain airspace configurations or situations may be limited. By incorporating TTC-detected encounters and their outcomes, the number of data points across the safety continuum increases, improving understanding of how aircraft enter particular trajectories and which controller and pilot practices contribute to more positive outcomes. This may ultimately inform formal rulemaking or the development of best practices to mitigate risk.

Over the next five to ten years, risk-based safety metrics may prove to be a valuable tool for managing operational risk. This raises important questions regarding their interaction with existing rules and standards. For example, could operations with well-managed risk, but which are not fully compliant, be considered acceptable?

Such scenarios are possible, and insights derived from TTC data and actualised outcomes may support innovative approaches to proposing new or modifying existing rules and standards. These insights could also inform best practices through analysis of TTC data combined with contextual information and advanced models (e.g. machine learning and artificial intelligence) to quantitatively support Safety Risk Management (SRM).

When regulatory focus is placed solely on efficiency and non-compliance, the inclusion of risk information helps balance safety and efficiency considerations. This balance supports informed discussion on how efficiencies may be gained while maintaining safety limits. The resulting data and insights could be used to establish new or modified standards that provide a more appropriate, specific, and defensible basis for future compliance requirements. What is currently defined as compliance could, over time, evolve into a more dynamic model based on real-time decision frameworks tailored to specific operational situations.

ANSPs have, for more than five years, successfully compiled and applied TTC results to optimize Air Traffic Management (ATM) while enhancing safety. International regulators may, in time, develop standards for the adoption and use of TTC, similar to existing compliance-based standards.

How would an ANSP use a TTC Algorithm?

The implementation of TTC algorithms and the availability of their outputs represent an initial step toward transitioning to a Risk-Based Safety Management (RBSM) approach. TTC algorithms identify interactions between aircraft with the potential for risk, regardless of whether those interactions are compliant. Depending on the size of the ANSP and the tuning of the algorithm, TTC analysis may generate thousands of potentially risky encounters per day. While this volume of data represents a significant asset, it also presents challenges for the adoption of RBSM methodologies, particularly in organisations where safety analysis has traditionally focused on individual non-compliant events.

To effectively use TTC data and determine both the level of operational risk and the factors influencing risk variation, ANSPs must establish downstream processes capable of analysing TTC outputs at an aggregate level. These processes may support a range of objectives, including (but not limited to):

- Monitoring the Day-to-Day Operation;
- Monitoring the Effectiveness of Change;
- Proactive and Efficient Airspace and Procedural Redesign;
- Analysis of Airspace Excursions;
- Probability of Collision, and;
- Barrier Analysis Review (BAR).

Monitoring the Day-to-Day Operation

TTC could be used to monitor the full extent of an ANSP's airspace and proactively identify and visualise hotspot areas where aircraft are frequently projected to come into close proximity with high rates of closure, irrespective of airspace rules, procedures, design, or the rules of the air being applied.

It is assumed that an ANSP is familiar with their own operation and that showing them areas of convergence that they are already aware of has little benefit or value, but is that completely true?

Capturing data and combining it with the ability to visualise these areas and analyse them alongside additional contextual information represents a significant step beyond intuitive “knowing” (i.e. suspecting). Access to structured data elevates understanding further. When risk assessments are driven solely by incidents of non-compliance, it may not be clear why certain aspects of operations result in higher-risk or more significantly non-compliant events.

While air traffic control (ATC) may conduct a procedure multiple times per day and expect these to appear as hotspots on a TTC dashboard, the ability to filter and curate encounters using additional contextual data allows for the identification of encounters that are largely similar, except for specific discriminating factors. Such contextual data may include:

- Route structure;
- Procedures;
- Runway configuration/geometry;
- Weather conditions;
- Unstable/high-energy approaches;
- Converging Runway Operations;
- Go-arounds;
- TCAS;
- Voice/text;
- Training, and;
- Fatigue.

By identifying individual hotspot areas, ANSPs can gain a deeper understanding of risk, operational context, and existing risk mitigations, whether procedural, related to airspace design, or linked to potential over-reliance on personnel. This understanding can support the identification of targeted remedies and mitigations. The resulting assessments may then be documented to inform Safety Risk Management (SRM), providing an objective picture of an ANSP's operational risk, associated risk control measures, and monitoring arrangements.

Monitoring the Effectiveness of Change

When operations are monitored on a regular basis, TTC tools can also be used to establish a risk baseline for each operation. This involves continually monitoring and trending the frequency of significant interactions across different sections of airspace to understand the full spectrum of risk and to associate like-to-like encounters. These encounters can then be further analysed to identify mitigating, aggravating, and other contextual factors influencing both outcomes and actual risk.

This data-rich approach supports the development of a statistically relevant operational baseline that can be used to monitor the effects of traffic growth on safety and efficiency. In addition, it can be used comparatively to project the risk associated with proposed changes and to validate the actual risk of implemented changes to airspace, procedures, equipment, services, and training. Such information supports more data-driven SRM via thorough risk assessment processes and effective mitigations, as well as Safety Assurance through relevant, measurable, and quantifiable safety performance metrics and thresholds.

Proactive and Efficient Airspace and Procedural Redesign

A well-established, data-driven operational baseline—supplemented by projections of future risk exposure based on anticipated traffic growth or other factors—provides a robust foundation for modelling. This enables objective identification of airspace areas and associated procedures that may benefit most from future redesign.

The combined use of TTC data and actualised outcomes in modelling current airspace, as well as simulating airspace and procedures before and after proposed changes, enables quantitative comparison of design alternatives. TTC data also supports objective evaluation of the safety of future designs by ensuring that any collateral hazards are identified, assessed, and effectively mitigated. This approach supports maximising the benefits of proposed changes while managing associated risks.

A notable opportunity lies in identifying areas of airspace that experience high traffic volumes but comparatively low frequencies of close-proximity encounters, to better understand why this is the case and to apply relevant lessons elsewhere.

Analysis of Airspace Excursions

Aircraft transitions into different airspace (e.g. Class B excursions) are closely monitored and reported as compliance issues. To better understand the actual risk associated with these events, analysis of a substantial set of TTC results—enriched with actualised outcomes and additional indicators—can provide a more accurate and specific picture of the risks involved.

By using risk as a discriminator to prioritise areas for further study, this data can be used to model excursions and resulting interactions, supporting the identification of changes to airspace design or procedures that tend to place aircraft on trajectories leading to such excursions.

Probability of Collision

Collision risk modelling is a valuable metric that often relies on compliance data and a limited number of serious incidents or accidents as indicators of actual risk. Since TTC represents a mathematical projection of the point of closest proximity and the time required for aircraft to reach that point, it appears to be a relevant indicator of collision probability when considered in aggregate.

Combining TTC outputs with the extensive data generated by such algorithms offers opportunities for further research. Initial work in this area includes the study “Probabilistic Analysis of Airborne Collision Risk by Data Driven Approach” by Jingshu Wu, Firdu Bati, Ilya Pistryakov, and Donald B. McNeill. The results of this work suggest that higher-priority TTC encounters tend to align with risk distributions identified by more comprehensive collision risk models. However, further research is needed, including the integration of actualised outcomes and contextual factors, to more accurately predict risk across different conditions, aircraft types, and airspace environments.

Barrier Analysis Review

The Barrier Analysis Review (BAR) process, currently under development by the FAA, assesses safety data incorporating both compliance-based and risk-based information to identify systemic risk.

The review captures aggravating, mitigating, and observed factors related to the operation under study, along with an evaluation of the effectiveness of available safety barriers. The primary objective of BAR is to communicate actionable information to operational stakeholders. Aggregate analysis produced through this process supports the identification of systemic risk within the National Airspace System.

Through both individual encounter reviews and aggregate data analysis, BAR informs personnel at one or more facilities where potential risks or safety issues may be present and identifies causal and contributory factors that either exacerbate or mitigate risk. While individual BARs focus on specific operations, BAR output data—such as factors affecting barriers and barrier effectiveness—can be leveraged at an aggregate level to assess the overall health of safety barriers across the system.

BAR is enhanced through the enrichment of TTC data with contextual metadata associated with each encounter, enabling more efficient identification of relevant operational scenarios. This enrichment also supports more accurate identification of causal factors, making BAR findings clearer, more actionable, and more meaningful for facilities undergoing review.

Safety Culture

The potential influence of applying a RBSM approach on safety culture within an ANSP remains uncertain, and several important questions merit consideration.

Would the implementation of TTC better highlight risks and thereby positively increase awareness of intrinsic risk?

It is possible that identifying risk areas where no past compliance anomalies have been observed could raise questions about whether compliance parameters alone are sufficient to manage collision risk effectively. In such cases, operational management might become aware of a particular “hot spot” and take appropriate action, for example by adjusting route structures during the next airspace design update cycle.

Conversely, there is also a possibility that the outcomes of TTC analysis could be discounted where no history of losses of separation exists and the operation is considered stable and operationally sound. This raises further questions:

Could an RBSM approach inadvertently distract operators and reduce focus on compliance, and if so, what new risks might be introduced?

There is little debate that RBSM aligns closely with the fundamental definition of safety as the management of risk to an acceptable level. However, questions arise if RBSM becomes the primary decision-making factor for operations. For example, could an ANSP experience repeated formal losses of separation because aircraft are permitted to descend earlier, without waiting for five or three nautical miles of separation, on the basis that aircraft trajectories are diverging and assessed as low risk? At a minimum, early agreement with the ANSP’s oversight authority would be necessary to accept such an approach, assuming losses of separation continue to be monitored.

There is also a question as to whether RBSM could unintentionally contribute to a more punitive culture if used to highlight situations where unnecessary or avoidable risks were taken, even where operations remained compliant.

Safety culture and RBSM are fundamentally linked in the pursuit of safer skies. Safety culture reflects the collective attitudes, beliefs, and behaviors shared by employees and management regarding safety within an organisation. A strong safety culture promotes proactive engagement with safety practices and encourages individuals to prioritise safety in daily activities. When an ANSP fosters a positive and robust safety culture, personnel are more likely to participate actively in risk assessments and adhere to safety measures, thereby enhancing the effectiveness of RBSM strategies and the execution of Safety Management System (SMS) processes.

As discussed earlier in this paper, RBSM provides a structured means of identifying and addressing hazards before they result in incidents. While compliance monitoring has been highly effective in achieving the high levels of safety currently experienced by global aviation passengers, compliance alone is not sufficient.

The relationship between safety culture and RBSM is evident in the way they reinforce one another. A positive safety culture supports effective RBSM by encouraging shared responsibility for safety and open identification of hazards. In turn, RBSM provides systematic tools and processes for addressing risk, which can help strengthen and sustain a positive safety culture. Together, they form a dynamic and resilient safety system that not only meets regulatory requirements but also evolves to address emerging risks and challenges.

In practice, compliance-focused processes have often led to event-driven approaches, where safety actions are triggered primarily by incidents or accidents. Such approaches can undermine the development of a robust safety culture by reinforcing reactive rather than proactive behavior. In some cases, this can foster a blame-oriented environment in which the focus is placed on individual accountability rather than on understanding and addressing underlying systemic issues.

A blame-oriented culture may discourage the reporting of near-misses or potential hazards due to fear of retribution, limiting open communication and transparency. Without a consistent flow of information on emerging risks and near-misses, organisations miss opportunities to learn and improve. This can lead to repeated incidents and further entrench reactive safety practices.

The adoption of RBSM, supported by the application of TTC, offers an opportunity to move away from these negative cultural dynamics by shifting attention from individual events to overall operational risk. By focusing less on individual fault and more on systemic risk, safety information can be shared more openly, encouraging proactive discussion and implementation of mitigation actions to address identified safety concerns.

The Big Questions

Several key questions remain as ANSPs continue to develop and integrate risk-based analytical models alongside existing compliance processes.

First, the authors of this CANSO Think Paper consider that the development and implementation of TTC algorithms and processes offer significant potential to improve aviation safety. Many ANSPs are already developing or using TTC to complement existing safety processes and have identified both direct safety applications and secondary uses, such as in airspace redesign projects.


The intention is not to eliminate compliance-based rules and processes. Rather, TTC offers additional opportunities to identify events and contributing factors that may not be visible through compliance-based analysis alone.

Second, the adoption of TTC processes will require policy changes for ANSPs choosing to implement them, and in some cases cultural change. Compliance rules have long been used as a proxy for risk and have served the industry well. The continued development of TTC algorithms and processes has the potential to advance safety analysis further and contribute to improved safety outcomes globally.

Training challenges may also arise. Existing compliance-based standards are relatively straightforward for controllers to apply during live operations. Incorporating risk-based considerations, such as TTC, may be more complex. Over time, however, TTC-informed analysis could support the development of more tailored separation concepts, although such evolution will take time.

ANSPs should expect questions and, at times, resistance to change. Nonetheless, the incorporation of TTC represents an evolution of safety processes made possible by advances in computing capability and new analytical approaches. Clear communication of the purpose, process, and benefits of TTC will be essential to secure leadership support.

Finally, and perhaps most importantly, interested ANSPs and individuals are encouraged to participate in continued discussion on TTC for education, collaboration, and shared learning. This represents a developing area, and there remains significant opportunity for further growth and collective refinement.



Credit: NATS

References

1. Benefits of Risk Monitoring and Proactive Safety in the Aviation Industry
2. CANSO Best Practices submission paper from NATS
3. Wu, J., Bati, F., Pistryakov, I., McNeill, D.B., Probabilistic Analysis of Airborne Collision Risk by Data Driven Approach
4. ICAO, Safety Management Manual (Doc 9859)
5. ICAO Doc 4444 (PANS-ATM)
6. ICAO Annex 11 – Air Traffic Services
7. ICAO Annex 19 – Safety Management
8. ICAO, Global Aviation Safety Plan (GASP)
9. FAA, UAS Registration Statistics, latest available year
10. FAA, FAA Aerospace Forecast Fiscal Years 2025-2045
11. EASA, Innovative Air Mobility / Annual Safety Review
12. International Civil Aviation Organization. (2024). *Moving beyond compliance measurement* (AN-Conf/14-WP/881). Fourteenth Air Navigation Conference, Montréal, Canada, 26 August–6 September 2024.
13. EUROCONTROL, STCA and ATM Safety Performance Reports

Authors

First Author – Job Brügggen, Safety Officer, Luchtverkeersleiding Nederland (LVNL), j.brugggen@lvnl.nl

Second Author – Dr. Jon I. Parker, Data Analytics Lead, MITRE, jjparker@mitre.org

Third Author – Sam Espig, Head of NATS Safety Management System, UK NATS, sam.espig@nats.co.uk

Fourth Author – Bruno Ochín, Safety Intelligence Manager, Nav Canada, bruno.ochin@navcanada.ca

Fifth Author – Donald McNeill, Safety Analysis Group Manager, Federal Aviation Administration, donald.mcneill@faa.gov

Sixth Author – Joe Bello, Computer Scientist, Federal Aviation Administration (FAA), joe.bello@faa.gov

Technical Annex A: Benefits of Risk Monitoring and Proactive Safety in the Aviation Industry

Purpose of this Annex

This annex complements the main CANSO Think Paper by illustrating, through practical examples, the benefits of risk monitoring and proactive safety management in air navigation. It provides additional context on how risk-based tools are being developed and applied by ANSPs to enhance safety performance, while reinforcing the paper’s central argument that risk-based approaches can effectively complement compliance-based safety management.

Overview

Air Navigation Service Providers (ANSPs) play a critical role in ensuring the safe and efficient flow of air traffic. In recent years, several ANSPs and EUROCONTROL have developed risk-detection and risk-monitoring software tools to strengthen their safety management capabilities. The Federal Aviation Administration (FAA), Luchtverkeersleiding Nederland (LVNL), National Air Traffic Services (NATS), and EUROCONTROL have each built similar technical systems that convert aircraft location data into preliminary risk assessments.

These organisations developed risk-based tools to advance shared safety objectives. This annex discusses the benefits of deploying such tools, outlines their inputs and outputs, explains how risk measurement complements rule compliance, and—most importantly—serves as an open invitation to collaborate on data-driven risk detection and risk analysis methods in pursuit of continually improving aviation safety.

Introduction

Recently, three Air Navigation Service Providers (ANSPs) and EUROCONTROL developed risk-detection and risk-monitoring software tools. The Federal Aviation Administration (FAA), Luchtverkeersleiding Nederland (LVNL), National Air Traffic Services (NATS), and EUROCONTROL built similar computer systems that convert time-series aircraft location data into preliminary risk assessments.

These ANSPs were motivated by similar operational needs and circumstances, and by a shared responsibility to apply professional risk-based safety management in order to keep the flying public safe.

Specifically, this annex:

- Discusses the benefits of deploying risk-based tools;
- Provides a definition of risk-based tools that captures the characteristics of all four existing systems;
- Discusses how rule compliance and risk measurement complement one another; and
- Issues an open invitation for collaboration on risk-based safety management to further improve aviation safety.

Value of Risk-Based Tools

The FAA, LVNL, NATS, and EUROCONTROL developed tools to collect risk-based measurements in support of data-driven aviation safety. The underlying premise of these efforts is that more accurate measurement of risk improves the ability to manage and further reduce risk in aviation. Risk-measuring algorithms and software (i.e. risk-based tools) enable new ways to detect, monitor, manage, and ultimately reduce aviation risk.

One fundamental benefit of risk-based tooling is that it changes the conversation around safety. Prior to the availability of such tools, discussions on aviation safety often focused primarily on rule compliance by pilots and controllers. The introduction of risk-based tools allows this conversation to shift from “*Was a rule violated?*” to “*How much risk was present?*” This shift moves the focus from rule compliance alone toward measuring and understanding real-world aviation risk. Placing measured risk at the centre of safety discussions prompts critical questions such as: *What factors contribute to elevated risk, and how can current operational risk be assessed?*

This does not diminish the importance of monitoring compliance with rules. Procedures and rules—such as those contained in controller handbooks—were established to ensure an acceptable level of risk and to support safe and consistent decision-making, particularly where real-time risk calculation is not feasible. While non-compliance does not automatically equate to elevated risk, it is often a contributing factor in significant safety events. ANSPs that prioritise compliance issues based on associated risk, using risk-based tools, can more effectively target safety mitigation efforts where they have the greatest impact.

ANSPs that systematically detect preliminary risk within their airspace, supported by additional analysis, experience two key benefits:

- the ability to focus limited staff resources on analysing and learning from higher-priority events identified through risk-based tools; and
- the ability to more clearly demonstrate when risk is well managed and when risk levels are unacceptable and require mitigation.

Over the longer term, ANSPs that continuously measure and monitor preliminary risk will be better prepared for future operational changes. As new entrants—such as unmanned aircraft systems (UAS), advanced air mobility, and electric vertical take-off and landing—introduce new demands on airspace, ANSPs equipped with automated risk-monitoring processes will be better positioned to assess the safety implications of these changes. Establishing processes that measure today's low levels of risk supports the maintenance of high safety standards in the future.

Direct risk measurement also provides a valuable complement to counting accidents or rule violations. Incident-based metrics require undesirable events to occur and become less informative as incident rates approach zero. While low incident counts demonstrate that airspace is safe, they do not necessarily indicate whether safety is improving. Risk-based metrics offer a means to communicate continuous safety improvement even in highly safe systems.

Risk-measuring techniques can also be applied to synthetic air traffic data. These tools are not limited to real-time operational data and can be used to assess proposed changes to airspace or procedures before implementation. They can also support human-in-the-loop simulations and training by providing immediate feedback on both compliance and risk.

Finally, adopting risk-based methods supports system-wide risk analysis, trend identification, anomaly detection, risk optimisation, and the application of artificial intelligence (AI) and machine learning (ML). Many of these techniques require large volumes of data; risk-based tools help generate the data needed to support them. Systematically quantifying risk therefore enables new analytical approaches and deeper understanding of aviation safety.



Definition of Risk-Based Tools

The risk-based tools discussed in this paper are algorithms and/or software that convert aircraft location data into preliminary risk assessments.

Input Data

Risk-based tools process a stream of aircraft location data that, in principle, could be expressed in table form as shown in Table 1; however, real-world production systems are unlikely to provide live location data in a table format such as this. Production systems will often disseminate data using a streaming data delivery solution (e.g., Apache Kafka or the FAA’s System Wide Information Management (SWIM) platform). Production systems may also disseminate location data encoded in a format like ASTERIX or Automatic Dependent Surveillance – Broadcast (ADS-B). Regardless of how the location data is encoded and delivered, the number of aircraft location records (e.g., rows in a hypothetical table) will be large. A live feed of aircraft location data typically updates an aircraft’s position record once every one to 12 seconds. At that rate, more than a million location records could be disseminated to track a day’s worth of air traffic at a medium sized airport.

Vehicle ID	Time Stamp	Latitude	Longitude	Altitude
AA123	2024-09-15T22:19:27	32.85676	-97.41115	35000
BAW456	2024-09-15T22:19:27	33.63176	-97.18339	1800
DLH789	2024-09-15T22:19:27	32.76422	-96.88218	11000
AA123	2024-09-15T22:19:30	32.85676	-97.41115	34900
BAW456	2024-09-15T22:19:30	33.63176	-97.18339	1800
DLH789	2024-09-15T22:19:30	32.76422	-96.88218	11000

Table 1 - Illustrative Example of Time-Series Location Data

Output

Risk-based tools can create two different types of preliminary risk assessments: aircraft pair risk and airspace risk. The FAA, LVNL, NATS, and EUROCONTROL have created tools (see Technical Annex B) that estimate aircraft pair risk. Tools that measure airspace risk have been proposed and prototyped but not yet fielded into operations.

Aircraft Pair Risk

Aircraft pair risk describes the risk two specific aircraft pose to one another. This type of risk assessment is derived from the proximity between the two aircraft and their relative trajectories. Tools that measure aircraft pair risk align well with the aviation industry's tradition of analysing aircraft separation, which is a core compliance concern.

Figure 1 depicts a notional example of an aircraft pair risk plot. The left side of the figure shows an overhead view of two converging aircraft: the bottom aircraft turns right, and the aircraft continue on parallel courses. The right side of the figure shows a notional plot of what the risk versus time graph would look like given these trajectories. Risk starts low because the aircraft are far apart. Risk grows rapidly as physical separation is shrinking in the presence of a converging flight path. When the bottom aircraft turns right, the risk measurement rapidly decreases because the physical separation is no longer decreasing. At this point the risk remains stable.

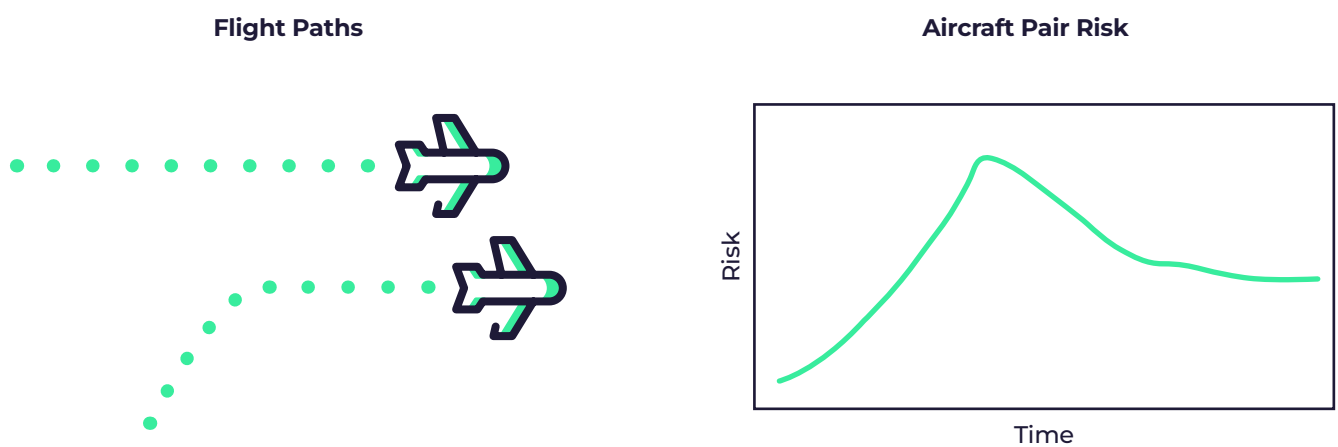


Figure 1 - Aircraft Pair Risk

We assert that an algorithm that measures aircraft pair risk will at a minimum have the following properties:

- **Positivity:** Aircraft pair risk is always a non-zero, positive number.
- **Reflects Physical Separation:** Aircraft pair risk decreases as the physical separation between two aircraft increases.
- **Reflects Time Separation:** Aircraft pair risk decreases as the closure rate between two aircraft decreases.

Algorithms that measure aircraft pair risk are sometimes referred to as Time to Conflict (TTC) tools, in that risk-based tools represent a potential estimate of “the time until two aircraft will come into conflict.”

Airspace Risk

The second type of preliminary risk assessment is airspace risk, which describes the risk in the entirety of the airspace. The value of extending from measuring simple pair risk to measuring total airspace risk is clear when you consider the question: What happens if a third aircraft is involved in an encounter? This question exposes a limitation of measuring only pair risk. When three or more aircraft are involved, encounter methods that only support pair-wise analysis may not properly capture the aggregate risk of the situation.

Measuring aircraft pair risk is simpler than measuring airspace risk. When measuring pair risk, there are only two aircraft to consider, not an unknown number of aircraft (n). When measuring pair risk, there is only one separation distance to track, not “n choose 2” separation distances. Figure 2 depicts an example of an airspace with many aircraft simultaneously aloft. Each green dot represents an aircraft, and each connecting line in the figure represents an aircraft-to-aircraft interaction that must be managed. We conjecture that methods to compute an instantaneous risk measure for an entire airspace could help us understand how total airspace risk changes as the number of airborne aircraft increases.

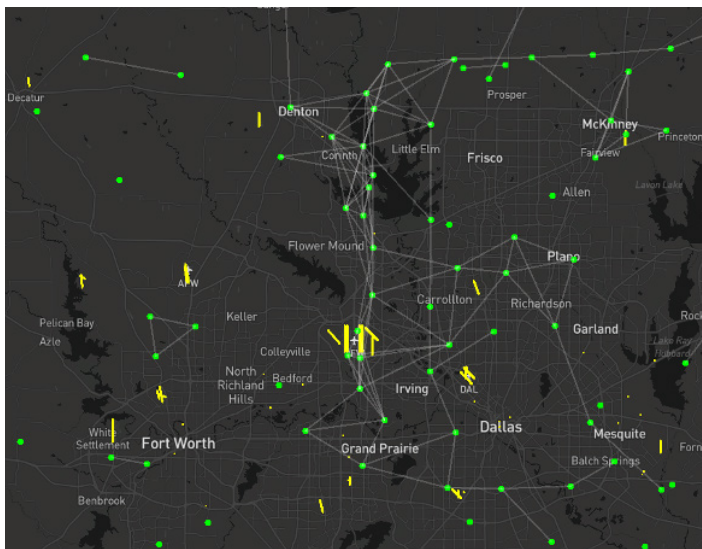


Figure 2 - Example of Numerous Aircraft-to-Aircraft Interactions in an Airspace

Limitations

Risk-based tools that process only aircraft location data will produce **preliminary** risk assessments. These preliminary risk assessments will remain useful and contain actionable information. However, they cannot be the definitive measure of risk because they are not fully informed and will not reflect crucial contextual information, that data for which is not available to the algorithm. The following are examples of contextual information that would influence a comprehensive risk assessment:

- Pilot intent
- Air traffic control intent
- Communication between air traffic controllers and pilots
- Weather
- Use of, or deviation from, an established arrival or departure procedure
- Airport runway configuration
- Accuracy and refresh rate of raw location data
- Existing controls and mitigations

Preliminary risk assessment methods that integrate these contextual data features could be designed and deployed, but context-aware methods will likely be built only after tools that measure preliminary aircraft pair risk are more widely available. Context-aware methods require more engineering effort to build and deploy because they must harness multiple disparate data feeds, including contributions by subject matter experts.

The Risk Assessment Score

Risk-based tools generate preliminary risk assessment scores; but what do these numeric risk scores mean? We can answer this question through an event-specific lens or through a statistical lens. For example, we may respond with, “The preliminary risk assessment score x was produced because lateral separation was a, vertical separation was b, and the closure rate was c.” Alternatively, we may answer, “The preliminary risk assessment score x was in the forty-fifth percentile of all observed risk assessment scores.”

Notice that neither of these responses focus on the actual value of x. One focuses on real world aspects of safety (i.e., physical separation and closure speeds), while the other focuses on the relative placement in a distribution.

The numeric risk assessment scores themselves are almost irrelevant. It does not matter if the score is 0.005, 5, or 50. The scores are an internal implementation detail.

Risk-based systems compute preliminary risk assessment scores to detect events with outlying characteristics and to characterise the risk distribution of all operations. Systematically computing preliminary risk assessment scores allows us to clearly define what “the risky 0.1% of all aircraft operations” means. We can then collect data on the subset of aircraft operations with the highest level of preliminary risk, study this subset, learn what real-world contextual factors may cause these events, and update operational aviation procedures to reduce these risk-factors.

Risk Complements Compliance

Many aspects of the aviation industry are governed by rules and regulations, which dictate pilot qualifications, air traffic control procedures, aircraft manufacturing, and more. These rules and regulations have clearly made a major contribution toward high safety standards, given the incredible safety level of today’s air travel. Rules and regulations provide pilots, air traffic controllers, and aircraft manufacturers with concrete structure to help everyone contribute to safe air travel daily. Ensuring compliance to aviation rules and regulations must remain a core pillar of the industry’s safety strategy because this rule driven structure has achieved excellent safety outcomes.

It is important to acknowledge that rule compliance and aviation risk are separate concepts. Measuring rule compliance does not directly measure the risks to which aircraft are exposed while in the operation. Figure 3 depicts the two possible options for how non-compliant aircraft operations and aircraft operations with higher levels of risk may overlap.

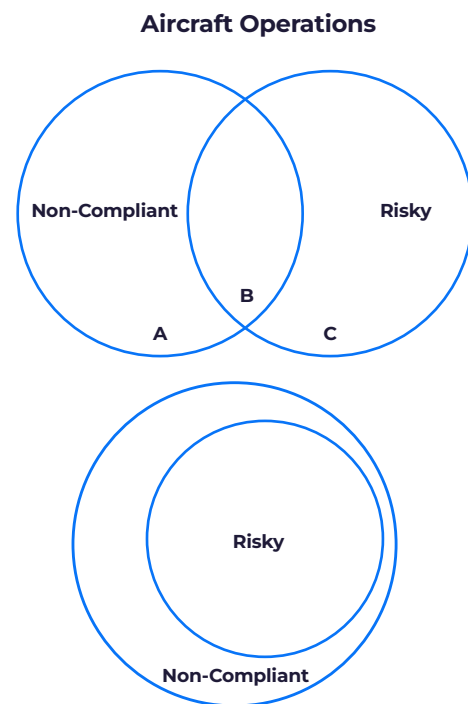


Figure 3 - How risky and non-compliant aircraft operations may overlap

The top part of Figure 3, non-compliant aircraft operations and risky aircraft operations do not fully overlap. The bottom part of Figure 3 illustrates a situation where all risky aircraft operations are non-compliant. Which part of the figure matches reality?

For the sake of argument, assume the top part of the figure is more accurate. In this case, what should the aviation industry do about the risky but compliant aircraft operations in region C of the figure? Are the safe yet non-compliant operations in region A overly constrained by rules and missing efficiency opportunities? How much do the non-compliant and risky regions overlap (i.e., how big are regions A, B, and C)?

Alternatively, assume the bottom part of Figure 3 is more accurate. In this case, how many rules (i.e., barriers) need to be broken before air travellers are exposed to too much risk?

How non-compliant does an individual aircraft operation need to be for it to enter the risky region?

Figure 4 depicts a situation at the boundary of compliance. By rule, two aircraft must maintain 3.0 nautical miles of lateral separation. Compliance to this rule is a binary variable. The separation requirement was either met or violated. As such, a plot of compliance versus time (see Figure 5) will result in a flat line or a step function, depending on whether the compliance boundary is crossed.



Figure 4 - The boundary of compliance with standard separation

Risk behaves differently from compliance. A small change to aircraft separation will yield a small change to risk. A corresponding plot of risk versus time, as shown in Figure 5, will display a smooth, continuous, differentiable function.

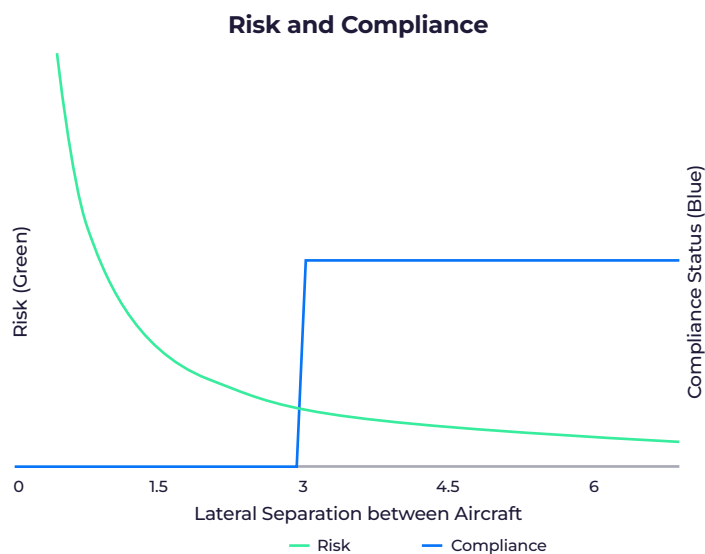


Figure 5 - Risk and compliance

This marks an important difference between compliance and risk. The derivative of compliance versus time holds virtually zero information. The slope is almost always zero and cannot support some ML methods. Alternatively, the derivative of risk versus time is a smooth function that supports more ML methods. If we develop tools to measure preliminary risk from moment to moment, then we open the door to detecting sudden changes in risk (i.e., the slope of the risk versus time function increases sharply). Knowing where the inflection points are in the risk versus time function will help us manage risk and understand the buffer zone between non-compliant and risky aircraft operations.

Gathering the hard data to answer the questions posed above will require measuring preliminary risk with the best available tools. While this attention to risk levels will complement the attention paid to compliance, it cannot replace it. Nonetheless, measuring preliminary risk will unlock new opportunities to improve safety.

Invitation to Collaborate

This paper is an open call to the aviation community to share, collaborate, and introduce ideas on how to measure preliminary risk using TTC tools and aircraft surveillance data. The Safety Intelligence Workgroup would like to seed the conversation with the questions below, but only as a starting point and with the expectation that new questions will be explored and consensus will emerge. Only by working collectively can the aviation community iterate to an industry standard on how to better assess preliminary risk using TTC. Some key points and questions to begin with are:

- What are the best tools and approaches for working with surveillance data?
- How can we best convert surveillance data into risk assessments?
- After collecting and analysing individual events, what is the best next step to aggregate data?
- How does contextual information like weather, airspace configuration, and pilot intent influence risk assessments?

This is an opportunity for all aviation professionals who have a passion for analytics and airspace safety to weigh in and build the next generation of Risk-based tools. Even if absolute consensus cannot be achieved, the conversation and collaboration will educate the community on what is possible and, from there, take away the best tools and ideas to better understand risk in the airspace.

OpenARIA


The FAA and MITRE Corporation's Center for Integrated Transportation built a TTC tool named Aviation Risk Identification and Assessment (ARIA). An open-source edition of ARIA (OpenARIA) has been publicly released on GitHub at: <https://github.com/mitre-public/open-aria>.

This work was open-sourced to facilitate building a community for anyone interested in viewing aviation through a risk-based lens. The repository's contents will make it significantly easier for others to work with aircraft location data.

The project contains:

- A full implementation of a risk-based tool,
- A fast solution for identifying aircraft pairs that warrant risk-based analysis,
- Tooling that simplifies working with streaming aircraft location data, and
- A mechanism for sharing benchmark datasets.

OpenARIA is available for immediate use, contribution, extension, and critique.



Technical Annex B: Shared Implementation Details

These four TTC tools have similar implementations. All four methods repeatedly recompute the *closest point of approach* (CPA) between both aircraft. Computing the CPA solves for:

- When the aircraft will arrive at their *closest point of approach* (i.e. minimum future physical separation).
- The lateral separation between the aircraft when the CPA occurs.
- The vertical separation between the aircraft when the CPA occurs.

Once the CPA is computed, it is used to assess the situation facing the two aircraft in the encounter. Table 2 lists the formula each ANSP uses to assess encounters. The FAA, UK NATS, and EUROCONTROL use formulae where lower values indicate lower physical separation and higher risk. The formula LVNL uses produces an integer between 0 and 15 where higher scores indicate higher risk.

ANSP	Formula
EUROCONTROL	$\frac{TimeToCPA}{30 \text{ sec}} + \frac{LatSep@CPA}{1 \text{ nm}} + \frac{VertSep@CPA}{250 \text{ ft}}$
FAA (and OpenARIA)	$\left(\frac{TimeToCPA}{30 \text{ sec}}\right)^2 + \sqrt{\left(\frac{LatSep@CPA}{0.25 \text{ nm}}\right)^{2.5} + \left(\frac{VertSep@CPA}{150 \text{ ft}}\right)^{2.5}}$
LVNL	$\begin{aligned} & \max\left(0, 5 - \text{floor}\left(\frac{TimeToCPA}{10 \text{ sec}}\right)\right) \\ & + \max\left(0, 5 - \text{floor}\left(\frac{LatSep@CPA}{0.5 \text{ nm}}\right)\right) \\ & + \max\left(0, 5 - \text{floor}\left(\frac{VertSep@CPA}{150 \text{ ft}}\right)\right) \end{aligned}$
UK NATS	$\frac{TimeToCPA}{20 \text{ sec}} + \frac{LatSep@CPA}{0.25 \text{ nm}} + \frac{VertSep@CPA}{150 \text{ ft}}$

Table 2 - Similar Formulae

There are also differences between the four tools and the input they receive. For instance, the FAA, LVNL, UK NATS, and EUROCONTROL methods do not receive aircraft location data from the same sources. Consequently, their input data will have different levels of accuracy. The aircraft location messages will also arrive at different frequencies. Update rates range from 1 hertz to 0.083 hertz. The quality of the input data may increase or decrease the value of sub-systems that reduce noise in either the input data (e.g., applying polynomial curve fitting to the latitude, longitude, and altitude measurements) or the output data (i.e., applying a moving average to the scores derived from Table 2's formulae).

