UNIVERSITY OF CAMBRIDGE | Aviation Impact Accelerator

Technical guardrails for contrails policy development

Principles, constraints, and tradeoffs

The Aviation Impact Accelerator
The Aviation Impact Accelerator is the global systems change initiative in aviation, based in the University of Cambridge. Combining data modelling and expertise across industry, governments and academia, we drive the transition to climate-neutral aviation.

The AIA's work draws upon the University of Cambridge's academic expertise and research rigour, together with the engineering capabilities of the Whittle Laboratory and the sustainability and policy expertise of the Cambridge Institute for Sustainability Leadership (CISL).

Authors
Peter Crookes, Dr Samuel Gabra, Jay Green, Noli Shelala, Ana Stojanovic

Acknowledgements
The authors would like to thank the significant number of individuals and organisations that have contributed their time and expertise to the development of the technical guardrails put forward in this report.

Citing this report
Aviation Impact Accelerator (2026). *Technical guardrails for contrails policy development - Principles, constraints, and tradeoffs*. Cambridge, UK: University of Cambridge.

Published
18 November 2026

Disclaimer
This document is published by the Aviation Impact Accelerator, an initiative led by the University of Cambridge. The Aviation Impact Accelerator aims to collect evidence and knowledge by engaging a wide range of stakeholders. The findings and conclusions expressed within do not necessarily represent the views of the University of Cambridge, nor those of the Aviation Impact Accelerator's partners, funders and collaborators.

Copyright
Copyright © 2026 University of Cambridge. Some rights reserved. The material featured in this publication is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International Licence (CC BY-NC-SA 4.0).

Table of contents

1 Executive summary 2

2 Framing the problem 4

2.1 Introduction 4

2.2 Policy criteria 6

3 Scientific foundation for contrails 8

3.1 Contrail formation and climate relevance 8

3.2 Sources of uncertainty 9

4 Mechanisms for contrail mitigation 11

4.1 Types of contrail mitigation mechanism 11

4.2 Navigational avoidance 13

4.3 Scheduling avoidance 19

4.4 Fuel composition 22

4.5 Aircraft design modifications 26

4.6 Comparing mitigation mechanisms 30

5 Technical constraints on navigational contrail avoidance 32

5.1 Constraint taxonomy 32

5.2 Observation infrastructure 33

5.3 Meteorological data 36

5.4 Decision support systems 38

5.5 Operational procedures 44

5.6 Regulation and standards 49

5.7 Summary 59

6 Policy scope: targeted or universal? 60

6.1 The efficiency argument 60

6.2 Taxonomy of targeting dimensions 60

6.3 Costs of regulatory targeting 61

6.4 Where targeting sits for navigational avoidance 63

6.5 Is targeting worth it? 63

6.6 Who carries which obligation 64

6.7 Summary 65

7 Conclusion and recommendations 66

7.1 Conclusion 66

7.2 Recommendations 66

7.3 Policy selection 68

7.4 Closing statement 71

8 Appendix 72

8.1 The flight-planning process 72

References 74

1 Executive summary

Persistent aircraft contrails are one of the largest and most immediately addressable contributors to aviation's climate impact. Current evidence suggests that contrail warming is comparable in magnitude to aviation's cumulative CO_2 impact, yet remains largely absent from climate policy frameworks. Without mitigation, contrails and aviation CO_2 together could consume a significant share of the remaining global temperature budget by 2050. Early intervention matters: delaying deployment substantially reduces the climate benefit available from contrail mitigation.

The central finding of this report is that **pre-tactical navigational contrail avoidance** should form the foundation of near-term contrail policy. This approach uses weather forecasts and flight-planning software to make small route or altitude adjustments before departure, allowing aircraft to avoid atmospheric regions where persistent warming contrails form. It can be implemented using existing aircraft, existing flight planning processes, and existing air traffic management infrastructure.

The policy opportunity is unusually attractive. Research consistently shows that **2-5 per cent of flights are responsible for 80-90 per cent of contrail warming**, meaning that targeted action can capture a disproportionate share of the climate benefit while affecting only a small fraction of flights. Additional fuel burn is typically small, and the warming avoided is estimated to exceed the warming caused by avoidance manoeuvres by around two orders of magnitude.

The report recommends nine core policy guardrails:

1. **Safety comes first**, ensuring that all contrail-avoidance actions remain subordinate to air traffic control requirements and operational safety considerations.
2. **Lead with pre-tactical navigational avoidance**, structured as a portfolio-based, performance-driven framework.
3. **Use portfolio-level monitoring and reporting**, recognising that individual contrail attribution remains technically constrained and that compliance frameworks should focus on aggregate performance.
4. **Avoid reliance on sustainable aviation fuel (SAF) or future aircraft technologies in the near term**, as scientific uncertainty, supply constraints, and attribution challenges currently prevent them from serving as credible primary policy mechanisms.
5. **Invest in infrastructure and weather data**, particularly humidity observations, weather forecasting, decision-support tools, and interoperability standards needed for operational deployment.
6. **Do not conflate CO_2 and non-CO_2 effects**, keeping contrail and CO_2 obligations separate so that action on one cannot offset inaction on the other, while allowing flight planning to weigh the two against each other.
7. **Name a single coordinating owner**, such as the Department for Transport (DfT) or the Civil Aviation Authority (CAA), because no single actor controls the chain from forecast to flown trajectory.
8. **Adopt adaptive, ICAO-compatible regulation**, beginning with voluntary or incentive-based measures and tightening obligations as the credibility of monitoring, reporting and verification (MRV) and operational capability develop.

9. **Apply to all operators equally, target by forecast**, so that forecasts and routing software, rather than route or airspace boundaries, select the highest-impact flights.

For policymakers, the most important conclusion is that contrail mitigation should not be viewed as a future technology challenge. The operational capability exists today to begin meaningful reductions in aviation's non-CO_2 climate impacts. While scientific uncertainty remains, available models already provide sufficient skill to support carefully designed pilot programmes, phased deployment, and adaptive regulation. Waiting for perfect prediction capability would likely delay action indefinitely without materially changing the core policy case.

The United Kingdom is particularly well positioned to lead internationally. Through institutions such as the Met Office, NATS, the CAA, and the Department for Transport, the UK already possesses many of the operational and governance capabilities required to develop, test, and scale contrail mitigation measures. A phased UK-led programme, as proposed in the accompanying white paper, *Searching for blue skies: Policy roadmap for navigational contrail avoidance* [1], and its policymaker summary, would deliver near-term climate benefits while helping establish the standards, evidence base, and regulatory models likely to shape future international policy.

2 Framing the problem

2.1 Introduction

2.1.1 Aircraft contrails contribute to global warming

Condensation trails ('contrails') are the white streaks that aircraft leave across the sky. They form when hot exhaust from an aircraft engine mixes with the cold air, causing water vapour to freeze into tiny ice crystals [2], [3].

Most contrails fade within seconds to minutes and have no meaningful effect on the climate [4]. However, in certain atmospheric conditions, contrails can persist for many hours, spreading into broad, hazy sheets that act like a thin blanket in the atmosphere and trap heat that would otherwise escape to space [5]. With nearly 40 million flights per year, these persistent contrails collectively produce an annual warming effect comparable in scale to all of aviation's carbon dioxide (CO_2) emissions combined: an unintended consequence of air travel that has grown in significance alongside the industry [6], [7].

2.1.2 A call to action

Unlike CO_2, contrail warming is not currently captured by any UK aviation climate target or carbon pricing mechanism. Accounting for these non-CO_2 effects in UK aviation policy would give a more complete picture of aviation's true climate impact, making contrail management a meaningful complement to existing CO_2 reduction commitments.

Without contrail mitigation, aviation is projected to contribute 0.054 °C from contrails and 0.040 °C from CO_2 by 2050. Together, that is 19 per cent of the remaining temperature budget between today and the +2 °C limit set in the Paris Agreement. If contrail avoidance is phased in between 2035 and 2045, up to half of this budget could be recovered. A ten-year delay, however, reduces this benefit to just 2 per cent. Importantly, the additional warming from extra CO_2 due to fuel-burn penalties is two orders of magnitude smaller than the warming avoided by reducing contrails. In summary, contrail avoidance represents a clear, near-term climate opportunity for aviation and swift, coordinated international action could make a meaningful difference [8].

The goal of contrail management policy is straightforward: prevent the fraction of flights that create these persistent, warming contrails from doing so.

2.1.3 Turbulence and contrails: a simple comparison

The operational framework for contrail management may already exist within the aviation sector today because of a similar phenomenon: *turbulence*. At their core, contrail formation and turbulence are both products of atmospheric states: transient, location-specific conditions that emerge from the physics of the upper troposphere [9]. Aviators can, in principle, navigate around both.

Turbulence avoidance is now a mature operational practice, embedded in the daily workflow of pilots and air traffic controllers. This practice can broadly be broken into three steps:

1. **Identification of turbulence regions**: The UK Met Office hosts the World Area Forecast Centre in London (WAFC London), one of only two global centres designated by the International Civil Aviation Organization (ICAO). It produces Significant Weather (SIGWX) charts and Clear Air Turbulence (CAT) diagnostics derived from high-resolution Numerical Weather Prediction (NWP) models, providing pilots with 3D upper-air hazard forecasts [10].
2. **Pre-tactical planning**: SIGWX forecasts are used by airlines and pilots for route planning before take-off. Flight plans are then submitted to the Integrated Initial Flight Plan Processing System (IFPS), operated by the European Organisation for the Safety of Air Navigation (EUROCONTROL). If accepted, the plan is distributed to the air navigation service providers along the route (NATS in the UK) [11].
3. **Reporting**: Pilots submit meteorological reports for events assessed as safety-relevant, such as severe turbulence. Significant meteorological information (SIGMETs) and airmen's meteorological information (AIRMETs) are issued by the Met Office as the designated Meteorological Watch Office; NATS controllers are then required to relay this hazard information to affected aircraft, enabling pilots to make tactical (in-flight) deviations [12].

The result is a closed loop: forecast products feed pre-flight route planning, in-flight reports from pilots feed back into safety decisions and air navigation service providers (ANSPs) act as the clearance broker translating meteorological intelligence into trajectory decisions. Contrail management requires a similar architecture, and the United Kingdom is well placed to build it. The Met Office operates one of the two WAFCs, giving it direct institutional responsibility for the global upper-atmosphere NWP products that any operational forecast system would build upon.

However, these systems are not one-to-one. More detailed differences are discussed later in the report, but there are clear distinctions between turbulence avoidance and contrail management which prevent immediate system adoption. A more fundamental operational distinction is vertical resolution: contrail-forming layers are often only hundreds of metres thick, imposing finer-grained altitude precision requirements than what turbulence avoidance currently demands [13]. Avoidance also carries different decision weights: turbulence avoidance is directly coupled to aircraft safety, whereas contrail avoidance is a climate mitigation measure, which has implications for how obligations and incentives must be structured. Yet, while the turbulence system took decades to build, a contrail management system could inherit its architecture relatively quickly because robust forecast and reporting infrastructure are already in place.

2.1.4 Report structure

The report focuses on the key scientific concepts, principal constraints and most relevant policy considerations needed to assess contrail management options.

Section 2.2 introduces the framework for assessing the viability of contrail mitigation policy, which must meet the following criteria: Safe, Effective and Feasible.

Section 3 establishes the scientific foundation: how and why contrails form and persist, the mechanisms by which persistent contrails warm the climate and the concentration of that warming in a small fraction of flights.

Section 4 assesses the four mechanisms for contrail mitigation in turn — navigational avoidance, scheduling avoidance, fuel composition and aircraft design modifications — evaluating each against the framework in Section 2.2. Navigational avoidance is identified as the most promising near-term mitigation strategy.

Section 5 expands on the technical constraints on navigational avoidance, classifying each as investment-dependent or structurally persistent.

Section 6 applies the constraints from Section 5 to determine the appropriate scope of a navigational avoidance policy.

Section 7 sets out our conclusions and recommendations, alongside a policy checklist to follow.

2.2 Policy criteria

2.2.1 What should contrail policy achieve?

Contrail management policy aims to reduce the net climate warming contribution of aircraft contrails. This report assesses contrail management policy against three core criteria: **Safe**, **Effective** and **Feasible**. A policy that fails any one of these tests is unlikely to succeed in practice, regardless of its theoretical merits.

Safe

Any contrail management policy must operate within the existing regulatory framework for aviation safety. Avoidance manoeuvres must be compatible with air traffic management procedures, and safety-critical instructions must always take precedence over climate-motivated routing changes. A policy that introduces new risks to aircraft or airspace is not likely to be accepted regardless of its climate benefits.

Effective

A policy is only worthwhile if it delivers a genuine reduction in contrail warming. An effective policy must satisfy four conditions simultaneously:

1. **Net reduction in warming** - warming reduction must be measurable and greater than the associated increase in emissions.
2. **High coverage** - policy scope must capture as much warming as possible.
3. **High compliance** - stakeholders must be incentivised to comply as much as possible.
4. **Low disruption** - minimise costs such as fuel burn penalties or delays/cancellations, and distribute them fairly.

Feasible

Good intentions are not enough if the technical, operational and governance infrastructure needed to implement a policy does not yet exist. Feasibility requires an honest assessment of what the science can currently deliver, what capabilities operators and regulators actually have and which constraints are temporary gaps that investment can close versus structural limits that policy must be designed around.

Safety sets the boundary within which any policy must operate - a policy that breaches this and compromises on safety will not be considered. Within that boundary, effectiveness defines the conditions that policy must satisfy, while feasibility measures whether the science, infrastructure and governance needed to deliver those conditions are available. Policies that fall in the intersection of effectiveness and feasibility are viable. Scientific advancements and dedicated investment can increase feasibility, but there are structurally persistent constraints to contrail management that impose a theoretical limit on the maximum feasible effectiveness. Section 5.1 classifies the constraints on contrail avoidance in these terms.

2.2.2 Why technical guardrails are needed

From light-touch subsidies to multi-sector market-based schemes, the options for contrail management policy can have a wide range of complexity. As this complexity increases, remaining true to the goal of reducing net climate warming can become muddled amid the detail. The above criteria should guide the selection, design and evaluation of all contrail-management measures, and should be measured throughout the policy process through modelling, flight trials and post-adoption monitoring.

3 Scientific foundation for contrails

3.1 Contrail formation and climate relevance

This section provides the scientific foundation for the policy analysis that follows. It covers contrail physics, climate forcing and prediction, focusing specifically on what policymakers need to understand to evaluate contrail-management strategies.

3.1.1 What contrails are and why most are climatically insignificant

A contrail is the white trail left by an aircraft engine. Hot, humid exhaust gas mixes with cold air, briefly condensing into ice crystals. Whether one forms at all depends on a thermodynamic relationship between exhaust temperature, ambient temperature and air pressure.

However, formation is not what matters primarily from a climate perspective. Most contrails disappear within seconds to minutes, sublimating harmlessly into the surrounding air, and contributing only a small warming effect.

What matters is whether a contrail persists, which depends on a different atmospheric condition: whether the air at cruise altitude contains sufficient water vapour such that the ice crystals grow faster than they shrink due to sublimation, a condition known as ice supersaturation. A contrail that forms here can spread and persist for hours, eventually broadening into a diffuse, cirrus-like layer indistinguishable from natural cloud.

3.1.2 How persistent contrails warm the planet

Persistent contrails act like a thin, man-made layer of high-altitude cloud. They intercept heat radiating upward from the Earth's surface that would otherwise escape to space. During the day, there is a partial counterweight, as contrails also reflect some incoming sunlight back to space. For young contrails in strong sunlight, this cooling effect can even surpass the warming effect. At night, there is no reflective offset, making nighttime contrails more warming. Despite this effect, contrails are assessed as net-warming overall throughout their lifetime. As a consequence, flights in the evening and overnight, including the majority of long-haul transatlantic operations, produce contrails with a higher per-flight warming impact than equivalent daytime flights through the same airspace [14].

Winter contrails at northern latitudes warm more. Low solar angles, long nights and cold temperatures make mid-to-high-latitude winter routes particularly effective at producing forcing [15]. Spreading contrails matter most. As contrails age and merge into diffuse cirrus sheets, their total climate impact grows. A contrail an hour old may cover hundreds of times more area than when it formed [16].

A standard measure of per-flight climate impact is energy forcing [17], [18]. This measures how much net heat (in joules) a contrail traps per metre of flight path, integrated over its full lifetime. This measure captures all of the above factors. Its estimated distribution across flights is highly skewed, as the following section describes.

3.1.3 Most of the problem comes from a small fraction of flights

Contrail climate forcing is heavily concentrated. Using estimates from the Contrail Cirrus Prediction (CoCiP) model, research finds that roughly 2–5 per cent of flights are responsible for 80-90 per cent of total contrail warming [19]. The rest produce contrails that either do not persist, persist briefly or form in conditions where the warming effect is low. This concentration has three primary causes:

1. **Not all airspace is equally risky.** Ice-supersaturated regions cluster at specific altitudes (typically 8–12 km, or flight levels 260–390), around jet streams and frontal weather systems and with strong seasonal and geographic patterns [20]. At any given moment, most airspace poses little contrail risk; however, a small fraction poses high risk.
2. **Busy routes overlap with risky airspace**. The North Atlantic tracks, US transcontinental corridors and intra-European routes all pass frequently through high-risk regions. The combination of high traffic density and high risk of warming contrails drives a disproportionate share of total forcing [21].
3. **Evening and overnight flights are overrepresented.** Because nighttime contrails carry more warming per kilometre, the concentration of long-haul departures in evening banks means a relatively small number of routes and time slots account for much of the total impact.

Together, these patterns suggest that a small identifiable subset of flights, routes and times could be targeted to capture most of the potential climate benefit. This concentration also means the climate benefit of contrail-management measures is not all-or-nothing: even partial compliance or targeting of only the highest-forcing flights could deliver a disproportionate share of the achievable benefit, making the case for intervention robust to less-than-full operational success.

3.2 Sources of uncertainty

Several sources of uncertainty affect the reliability of energy forcing estimates, the effectiveness of contrail-management interventions and therefore the predicted effectiveness of any policy:

→ **Contrail modelling uncertainty** — Contrail models generally fall into two categories: physics-based and machine learning. The most widely used physics-based model is CoCiP. While CoCiP is simple and requires low computational load, this simplicity can introduce substantial errors, so it is important to note that the energy forcing estimates from CoCiP are just a proxy for the true climate impact of contrails[22], [23]. Machine learning models are a promising, but nascent alternative in the modelling space. These models require extensive labelled datasets to effectively learn and are still being tested with live flight trials and geostationary satellites [24]. Both modelling approaches rely heavily on the quality of their inputs for accuracy.

→ **Atmospheric modelling uncertainty** — With current weather forecast models, ice-supersaturated regions are difficult to locate precisely, which contributes to uncertainty in the effectiveness of avoidance measures [25]. In addition, the most widely used models, such as CoCiP, calculate energy forcing estimates which rely on accurate meteorological data. Errors in humidity, temperature and wind fields propagate directly into predicted contrail lifetime and radiative forcing. These constraints are

further discussed in Section 5. Atmospheric uncertainty is less central to fuel and aircraft-design measures, but it still affects the assessment of their climate benefit.

- → **Compliance and implementation rates** — Regions where stakeholder compliance is likely to be low can be identified using information such as air traffic density, airspace complexity and the complexity of region-specific regulatory and operational systems. These factors are particularly relevant to operational (navigational and scheduling) avoidance. For fuel and aircraft measures, uncertainty instead arises from supply-chain availability, certification, procurement decisions and fleet replacement. Estimates of compliance and implementation rates are highly uncertain, which in turn contributes to uncertainty in policy effectiveness.
- → **Limited observational records** — Contrail climate models are calibrated and validated against a geographically uneven observational record, drawn mainly from satellite imagery, a small number of dedicated flight campaigns and ground- or aircraft-based lidar measurements. Satellite detection is generally limited to contrails that are wide, persistent and optically thick enough to be resolved, while flight campaigns cluster around specific regions and traffic corridors rather than providing global coverage. This limits confidence in how well model performance generalises across regions. Additionally, Automatic Dependent Surveillance–Broadcast (ADS-B) coverage is not uniform across all airspace or altitudes [26], and missing or truncated trajectories may undercount contrail-forming segments.

4 Mechanisms for contrail mitigation

4.1 Types of contrail mitigation mechanism

Contrail mitigation mechanisms describe how contrail mitigation is physically achieved and are distinct from the policy instruments (mandates, incentives, voluntary schemes) that influence their adoption. This distinction matters because the same mechanism can in principle be deployed under different policy frameworks, and the feasibility of each mechanism must be assessed on its own technical and operational merits before considering how it might be regulated or incentivised.

4.1.1 Where mitigation fits in flight operations

Each mechanism acts at a particular point in the flight-planning process, and that point decides who makes the decision and how much atmospheric information is available when they do. The process has four phases, shown in Figure 1 and described in more detail in Section 8:

1. **Strategic planning** (months ahead): Airlines and slot coordinators fix which flights operate and when, but not their routes.
2. **Pre-tactical planning** (12–24 hours ahead): Dispatchers build and file the flight plan. The Network Manager (NM) reviews it and, just before departure, issues take-off clearance with a short departure window.
3. **Tactical management** (real time): Crew and ATC agree any deviations from the filed plan in flight.
4. **Post-flight review**: Airlines compare the filed plan with the route flown and record fuel burn, delays and deviations.

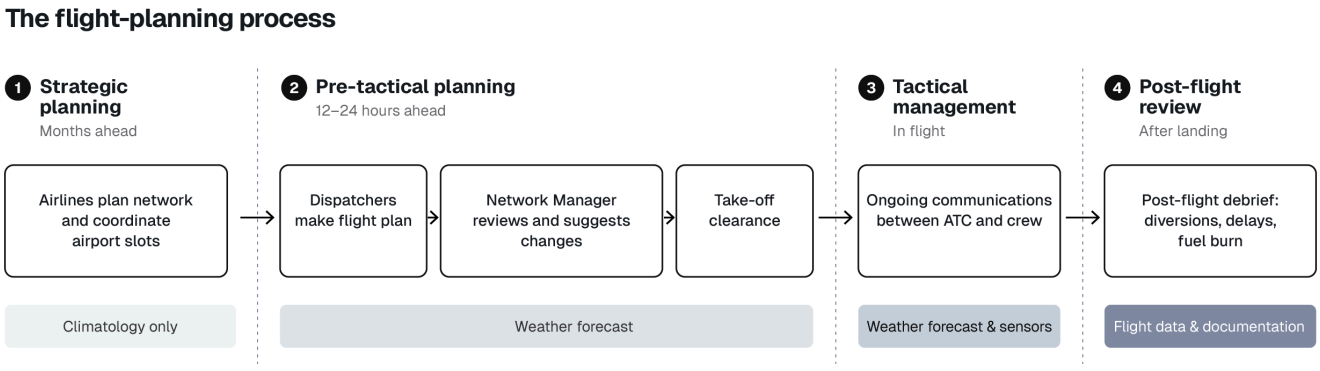


Figure 1: A simplified view of the typical flight-planning process.

4.1.2 Classifying mitigation mechanisms

Options for contrail mitigation fall into two broad categories according to their physical mode of action. These can be seen in Figure 2. Avoidance strategies prevent persistent contrail formation entirely by routing or timing flights so that they do not pass through ice-supersaturated regions (ISSRs). Property modification strategies instead allow flights to pass through ISSRs but alter the exhaust composition so that any contrail that does form has different microphysical properties minimising its radiative forcing. The distinction matters for policy design: avoidance strategies depend on atmospheric

prediction and trajectory flexibility, while property modification strategies depend on fuel supply chains, fleet renewal and aircraft manufacturing.

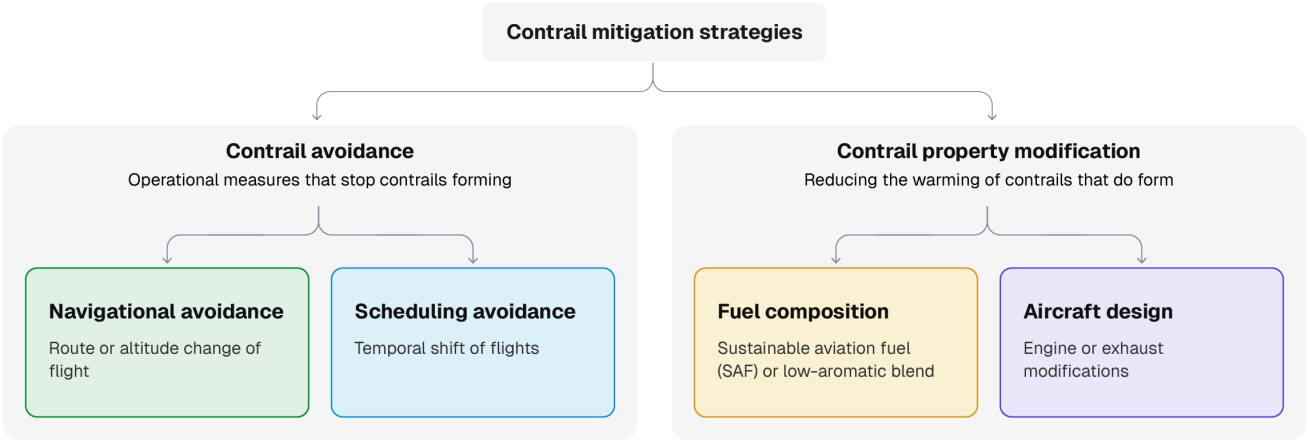


Figure 2: The classification of contrail mitigation strategies

The main mechanisms for contrail mitigation are [27]:

1. **Navigational avoidance**: This involves altering the flight path to avoid regions of the atmosphere that are conducive to contrail formation. This can include changes in altitude or route. This is the most operationally mature category and the primary focus of most current research and trials [28].
2. **Scheduling avoidance**: This involves adjusting the timing of flights to avoid periods when contrail formation is more likely. This can include delaying or advancing departures, or rescheduling flights to different times of day. Scheduling avoidance is less mature than navigational avoidance and may be more challenging to implement due to operational constraints and passenger demand.
3. **Fuel composition**: This involves using modified or alternative fuel as a means to reduce contrail formation. Sustainable aviation fuel (SAF) usage is a medium-term strategy that can be implemented more quickly than aircraft design modifications, but it still faces challenges related to production capacity, cost and supply chain logistics [29].
4. **Aircraft design modifications**: This includes changes to the aircraft's design, such as modifications to the engine or exhaust system, to reduce the likelihood of contrail formation. Long-term research and development is needed to bring these technologies to market, and they are not expected to be widely available in the near future [27].

Each mechanism intervenes at a different phase of the flight-planning process. Figure 3 adds these intervention points to Figure 1. Static scheduling restrictions and aircraft design modifications must be fixed at the strategic planning stage, since airlines commit to network schedules, slot requests and fleet assignments months or years in advance. Neither requires real-time atmospheric data, which allows them to be set early but makes them unresponsive to day-to-day atmospheric conditions. Pre-tactical navigational avoidance acts when dispatchers build the flight plan, typically 12–24 hours before departure, while dynamic scheduling avoidance is limited to the narrow take-off window assigned by ATC. Tactical navigational avoidance takes place in flight, and post-flight review supplies the data needed for monitoring, reporting and verification.

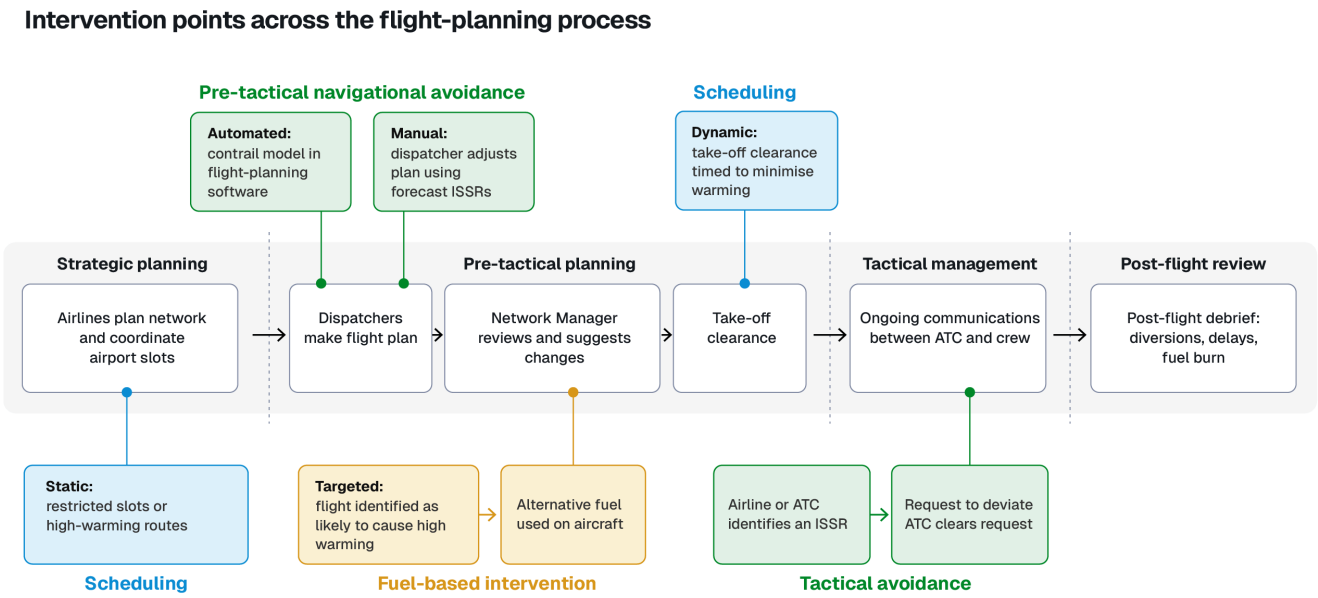


Figure 3: Intervention points for each mitigation mechanism, added to the flight-planning process in Figure 1.

The following sections describe the different types of mitigation mechanisms in detail, including the physical basis for the intervention, evidence on effectiveness and key implementation challenges.

4.2 Navigational avoidance

Navigational avoidance is the modification of a flight's trajectory by altering its route or altitude to minimise or prevent contrail formation. The two primary implementation modes are pre-tactical and tactical. Pre-tactical avoidance involves adjusting the flight plan before departure, while tactical avoidance involves making adjustments during flight based on real-time atmospheric conditions. This section explores the effectiveness of navigational avoidance in reducing contrail formation, the challenges associated with its implementation and potential strategies for optimising its use in aviation operations [30], [31].

This strategy relies on the shape of ISSR regions. Their pancake-like structure allows for targeted adjustments to flight paths to avoid contrail formation. In most cases, this would involve temporarily adjusting the altitude by two flight levels (2,000 ft). If the location of these regions can be reliably identified, navigational avoidance can be integrated into existing flight-planning and air traffic control systems. However, its effectiveness is highly dependent on the accuracy of weather forecasts and the ability to predict contrail formation, which can be challenging due to the complex interactions between atmospheric conditions and aircraft emissions.

Figure 4 shows the altitude change typically required to avoid a contrail-forming region. Because ice-supersaturated regions are relatively thin vertically, most cases can be avoided by shifting the aircraft by only one or two flight levels. In practice, this makes navigational avoidance a relatively simple operational intervention when the affected region can be identified in advance.

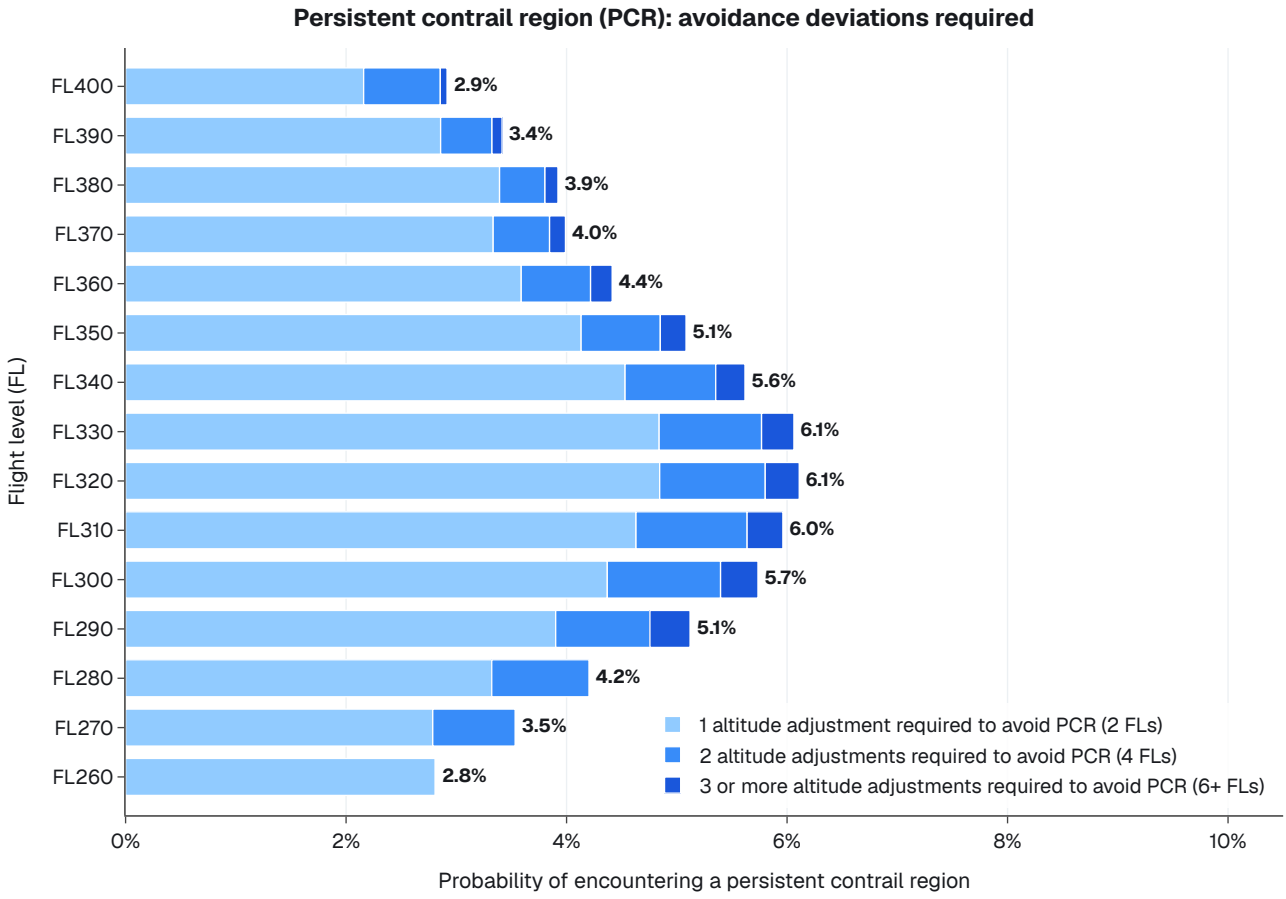


Figure 4: Most flights can avoid a contrail-forming region with only a small change in altitude. Source: Contrails.org, 'Contrail Regions 101: Part I'.

4.2.1 Pre-tactical avoidance

Pre-tactical decisions are those that occur before a flight. They are made by airline dispatchers during flight planning in coordination with the Network Manager. For contrail management, effective pre-tactical decisions primarily rely on the accuracy of weather forecasts and the ability to predict whether a warming contrail will form.

Avoidance at this stage can involve either a manual adjustment, in which the dispatcher modifies the planned altitude or route based on a forecast ISSR, or an automated adjustment, in which a contrail model embedded in flight-planning software generates a contrail-aware, optimised trajectory. Both require a usable forecast at planning time.

Meteorological forecasts generally become more accurate as their lead time shortens, so pre-tactical avoidance is more effective when implemented closer to departure. Flight plans are commonly developed a day in advance, which constrains the accuracy of pre-tactical avoidance without causing disruption [32]. However, even with accurate forecasts, pre-tactical avoidance may not always be feasible because of ATC constraints, fuel-efficiency considerations and other operational factors. Figure 5 illustrates how forecast accuracy affects climate-optimised flight trajectories generated using CoCiP. The source estimates that optimising a trajectory a day in advance reduces the accuracy of the contrail-avoidance plan by 15 per cent compared with optimising it immediately before departure.

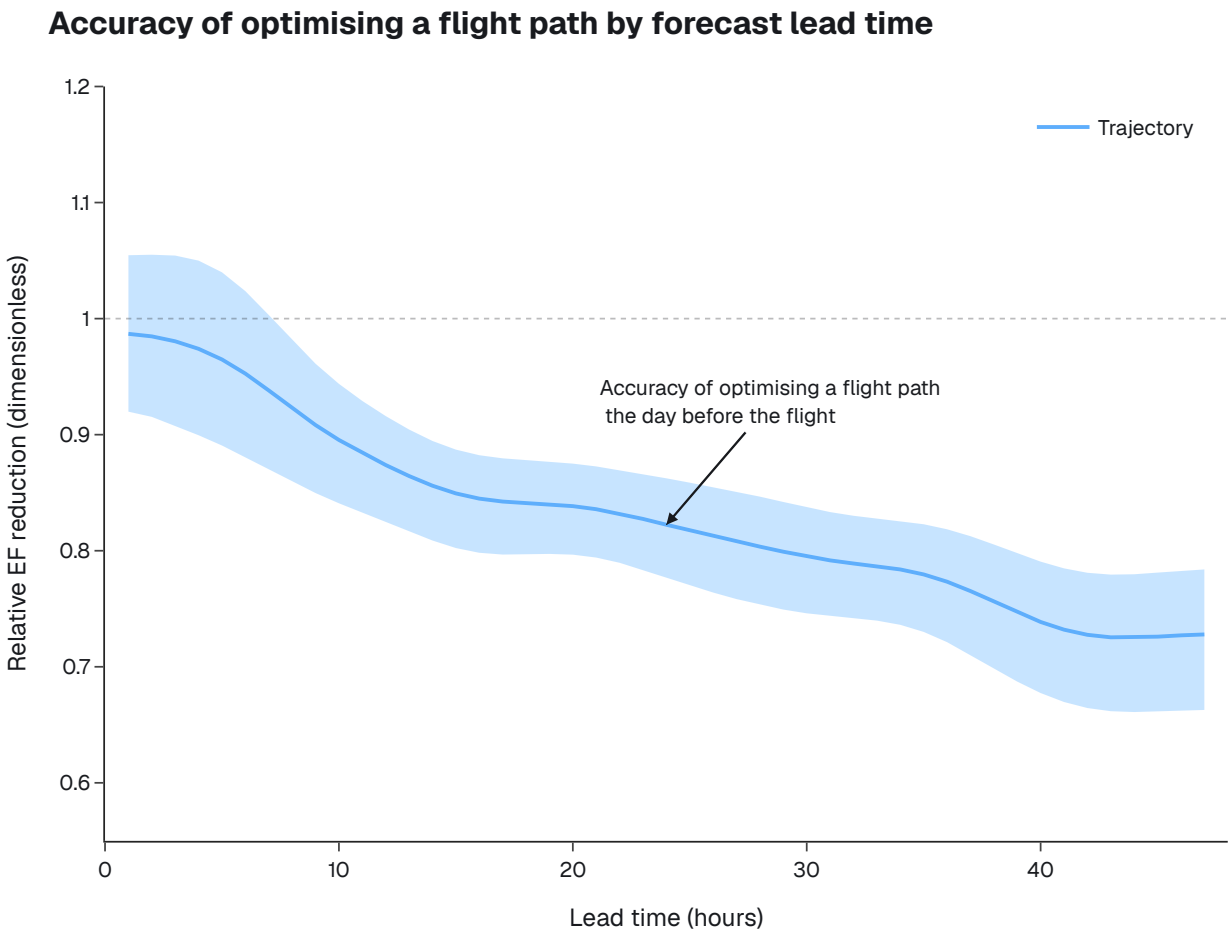


Figure 5: Simplified diagram showing how the accuracy of weather predictions affects the accuracy of a climate-optimised flight plan. Source: Data from Dean et al. 2025, 'Impact of Forecast Stability on Navigational Contrail Avoidance'.

Pre-tactical avoidance reduces workload for pilots and air traffic controllers compared with tactical avoidance because the flight path is planned and agreed upon before departure, reducing the likelihood of in-flight adjustments.

The main constraints of pre-tactical avoidance include the need for accurate weather forecasts, potential conflicts with air traffic control and the possibility of increased fuel consumption due to longer routes. Additionally, pre-tactical avoidance may not be effective in rapidly changing weather conditions, which can lead to unexpected persistence or dispersion after formation [33].

4.2.2 Tactical avoidance

Tactical avoidance involves adjusting the flight path during flight based on observed or forecast atmospheric conditions. This approach allows for more flexibility in responding to changing weather patterns and may be more effective in avoiding contrail formation. However, it also presents several challenges, including increased workload for pilots and air traffic controllers, as well as the need for real-time atmospheric data [34].

Tactical avoidance requires either an airline's operations team or air traffic controllers to monitor weather conditions and initiate a request to adjust the flight based on an observed or forecast ISSR region. These are referred to as 'airline-led' and 'ANSP-led' tactical avoidance, respectively. Both

approaches rely on established communication channels and procedures for adjusting flight paths. Third parties, such as a meteorological service, may provide or relay this information, while the flight crew and relevant air traffic controllers retain authority to agree to or decline a manoeuvre on safety and operational grounds.

The core requirements and challenges are similar in both ANSP-led and airline-led tactical avoidance. In either case, the controller and flight crew must agree on an alternative route, and either party can decline the request if the manoeuvre is judged unsafe or operationally infeasible, for example, because of insufficient airspace capacity. The key distinction is who identifies the contrail-forming region and initiates the deviation: in ANSP-led avoidance, the controller identifies the ice-supersaturated region and instructs the aircraft accordingly; in airline-led avoidance, the airline's operations team identifies the region from available forecast or onboard data and coordinates a deviation request with the flight crew. In both approaches, the deviation remains subject to agreement between the controller and flight crew.

The choice between these approaches depends on available resources and the objective of the contrail mitigation measure. Where the goal is to minimise contrail warming across a defined volume of airspace, ANSP-led avoidance offers greater control over which aircraft are diverted, helping to limit 'saturation' effects, in which the climate benefit of one diversion is eroded by other flights forming contrails in the same region shortly afterwards [35]. The controller's wider view of surrounding traffic also allows conflicts to be anticipated in advance, minimising disruption in the airspace. An airline-led approach, by contrast, is generally better suited to minimising the total warming impact of an individual flight, since it enables the airline's operations team to coordinate an alternative route based on conditions observed along the flight's trajectory, regardless of the airspace or ANSP under which the flight is operating. This distinction matters when warming is assessed on a per-flight basis.

4.2.3 The cost of navigational avoidance

For an individual flight, navigational avoidance generally requires only limited changes to existing operations. Scaling these interventions across the airspace is more challenging. Procedures that are practical during small-scale trials may become operationally infeasible at higher traffic volumes, particularly where deviations compete for limited airspace capacity. This creates a strong case for automating as much of the process as possible, including the identification of suitable deviations, coordination with air traffic control and the prioritisation of flights for intervention.

Reported trials provide important operational evidence on the implications of implementing navigational avoidance at scale. They indicate that, as traffic levels increase, reserving airspace for avoidance manoeuvres can become impractical. The Operation Blue Skies trial [36], scheduled for 2026–2028, will provide evidence on the feasibility of scaling up this mitigation measure over oceanic airspace. Further evidence on capacity constraints in the UK may need to be obtained through additional trials over busy domestic airspace. As aviation demand grows and congestion increases, airspace capacity may become the principal factor limiting the effectiveness of navigational avoidance.

Contrail-forming regions occur frequently enough that avoidance will need to be treated as a routine operational process rather than an exceptional intervention. Approximately 30 per cent of flights may form persistent contrails; simulations suggest that contrails may form over 5 per cent of flight distance [21], while contrail-forming regions cover 6 per cent of the airspace on average [37]. These estimates are informative but do not translate directly into a uniform operational requirement. The number of flights requiring a deviation will depend on the objective of the intervention, the geographical scope of the scheme and the atmospheric conditions at a particular time. In some circumstances, most flights may require a deviation, whereas in others no intervention may be necessary.

As understanding of contrail-forming regions has improved and flight-planning software has advanced, estimates of the fleet-wide fuel-burn penalty have decreased significantly. A conservative estimate places the penalty at around 2 per cent across the fleet, although some recent trials found no statistically significant increase in fuel burn [38]. The penalty is, however, unlikely to be distributed evenly across the fleet. The fleet-average fuel-burn penalty is therefore a useful summary metric, but any policy framework must also account for the possibility that the operational burden will fall disproportionately on particular operators, routes and areas of airspace.

Policy guardrail · Lead with pre-tactical navigational avoidance

Navigational avoidance policy should prioritise pre-tactical avoidance, with tactical avoidance used as a complement where controller workload and airspace capacity allow. Planning deviations before departure keeps avoidance within established flight-planning review processes, limits real-time workload and allows conflicts with other traffic to be resolved in advance. To remain practical at scale, the identification of suitable deviations, coordination with air traffic control and the prioritisation of the highest-warming flights should be automated within flight-planning and air traffic management (ATM) decision-support tools rather than left to manual judgement. Controllers and flight crews must retain the authority to decline any deviation on safety or capacity grounds.

4.2.4 Effectiveness of navigational avoidance

Navigational avoidance is the most operationally mature contrail mitigation strategy and does not require an overhaul of aircraft technology, supply chains or fuel certification. It can be implemented with existing aircraft and ATM infrastructure, requiring software updates to flight-planning systems rather than new hardware or fuels, so deployment can begin without waiting for aircraft or fuel changes [39].

Because contrail warming is concentrated among a small fraction of flights, targeting those flights for deviation can deliver substantial warming reductions while using resources more efficiently than broad measures applied across the entire fleet.

In ANSP-led avoidance, deviations can be managed centrally, conflicts with other traffic can be assessed systematically and the outcomes of avoidance decisions can be logged and audited against a consistent standard. The UK has more direct control over flight paths within NATS airspace, which could make it a more feasible candidate for implementing navigational avoidance strategies compared with regions with more fragmented airspace management.

4.2.5 Assessment against the policy criteria

Measured against the Safe, Effective and Feasible criteria (see Section 2.2):

Assessment

- → **Safe**: Altitude and route changes are routine manoeuvres that are already cleared through existing ATC procedures. Either the controller or the crew can decline a deviation on safety or capacity grounds, so avoidance remains subordinate to safety-critical instructions. Tactical avoidance adds controller and crew workload, which must be managed as traffic density increases.
- → **Effective**: Because contrail warming is concentrated in a small fraction of flights, deviating relatively few flights can deliver a large net reduction in warming. The climate impact of the additional fuel-burn emissions is estimated as two orders of magnitude smaller than the warming avoided [8]. Effectiveness is limited by forecast accuracy, which favours planning closer to departure, by saturation effects in busy airspace and by the difficulty of verifying avoidance for individual flights.
- → **Feasible**: Navigational avoidance uses existing aircraft, flight-planning processes and ATM infrastructure, so deployment does not depend on new hardware or fuels. Contrail-aware flight-planning tools exist but are not yet widely deployed, and airspace capacity may become the principal limit at scale. Tool deployment is investment-dependent, and capacity can be eased through ATM modernisation, but safety primacy and finite airspace remain persistent constraints.

Near-term role: primary. Pre-tactical navigational avoidance is the strongest near-term contrail mitigation mechanism, with tactical avoidance as a complement where controller workload and airspace capacity allow. Its practical limits, set by forecast skill, tooling and airspace capacity, are examined in Section 5.

4.3 Scheduling avoidance

Scheduling avoidance reduces the contrail warming effect by modifying the timing of flights rather than their trajectories. An implementation of this approach could involve dynamically adjusting flight schedules based on real-time atmospheric conditions. Taken further, scheduling avoidance could be implemented as a blanket restriction or curfew in order to utilise current understanding of contrail warming and its dependencies. This section explores the effectiveness of scheduling avoidance in reducing contrail formation, the challenges associated with its implementation and how it compares to navigational avoidance strategies.

4.3.1 Dynamic scheduling avoidance

For dynamic scheduling avoidance, where individual flights are shifted temporally to avoid specific atmospheric conditions, the effectiveness of the strategy depends on the length of the tolerable delay window. In this strategy, airlines or other stakeholders would need to monitor atmospheric conditions in real-time and adjust flight schedules accordingly. The length of the delay window would determine how much flexibility airlines have in rescheduling flights to avoid contrail formation. The foundation of this suggestion is that contrail formation is highly sensitive to atmospheric conditions, and that small changes in timing could significantly reduce the likelihood of contrail formation.

Within the existing flight-planning process, the natural intervention point for dynamic scheduling avoidance is take-off clearance. However, the calculated take-off time window for a flight (of around fifteen minutes) provides a very narrow opportunity for timing adjustments without cascading effects, and ISSR regions are sufficiently persistent at this timeframe that delays would rarely achieve meaningful avoidance. For a longer scheduling shift, the Network Manager and airline might agree to shift a flight's departure time, which could have cascading network effects. In a scenario where only one flight requests a scheduling change, the effects may be minimal and similar to typical operations. However, the effect of many airlines avoiding the same take-off times is unknown and could be significant.

Larger delay windows can have a substantial operational effect, as they may require rescheduling of multiple flights and coordination between airlines, air traffic control and airport authorities. This leads to increased costs for airlines and passengers, alongside potential disruptions to travel plans [40]. The operational success of this mechanism would depend on minimising this disruption and the time scale over which the flight schedule is adjusted. Delays of the order of hours compared to the original schedule may require airlines to pay for additional crew time as well as delay compensation.

Currently, there is no evidence to support the claim that dynamic scheduling avoidance is beneficial or even feasible within the current operational framework. The flexibility required for this method to achieve a meaningful warming reduction does not exist within the current framework since larger delays can have significant network effects.

4.3.2 Static scheduling avoidance

A more static implementation of scheduling avoidance involves incorporating restrictions on flying at high-warming times. An example of this would be targeting nighttime flights. Nighttime flights are known to have a stronger warming effect: during the day, contrails “block” incoming solar radiation but also prevent outgoing heat from escaping. At night, only the insulation effect remains and net radiative forcing is more positive [41]. In static scheduling avoidance, most nighttime flights would consequently shift towards daytime, minimising the warming impact. In one form, this shift could occur through further restricting night movement quotas already established at Heathrow, Gatwick and Stansted airports [42]. Using this approach, rescheduling all nighttime flights to daytime would therefore reduce the overall warming contribution of contrails without requiring detailed per-flight analysis. This approach is a more restrictive version of dynamic scheduling avoidance since it affects all flights that meet the criteria regardless of warming contribution. Static restrictions of this kind must be embedded at the strategic planning stage, since airlines need to know their operating constraints before they commit to network schedules, slot requests and fleet assignments.

Critically, the difference in warming between nighttime and daytime contrails is well established in contrail science, meaning this intervention does not depend on uncertain forecasting or atmospheric modelling. Stuber and Forster found that the global distance travelled by aircraft at night was less than 40 per cent, yet contributed to a disproportionate 60 per cent of the global annual mean forcing [43]. However, using general trends instead of per-flight data would mean that many flights which do not form contrails at night would be unnecessarily rescheduled, leading to increased operational disruption and costs. This approach would require significant changes to airline operations and passenger schedules, and may not be feasible for all routes. Nighttime airspace is also less congested, meaning shifting flights to daytime could increase delays, raise fuel burn and thereby increase CO_2 emissions — partially offsetting any contrail-related benefit. This trade-off would need to be accounted for in any cost–benefit analysis.

4.3.3 Effectiveness and feasibility

Scheduling avoidance can be implemented at a policy level as a structural intervention without requiring real-time decision-making by pilots or air traffic controllers. A blanket policy such as a night flight curfew does not require the atmospheric precision or continuous monitoring needed for navigational avoidance. However, scheduling avoidance has significant implications for airline operations, including disruptions to flight schedules, increased costs and reduced flexibility for passengers [44]. The more dynamic implementation, with short notice adjustments to flight schedules, is not operationally feasible for airlines while at the same time requiring monitoring of atmospheric conditions and coordination between multiple stakeholders [45].

ISSR regions move and evolve on the timescale of hours. Constraining the problem to only changing timing rather than readjusting the flight path may not be sufficient to avoid contrail formation in a reasonable way. Additionally, rescheduling flights at airports requires coordination between multiple stakeholders, including airlines, air traffic control and airport authorities. Air traffic control at airports is

constrained by the number of available runways and gates, as well as the need to maintain safe separation between aircraft. This can limit the ability to adjust flight schedules in response to real-time atmospheric conditions, particularly during peak travel times when airport capacity is already stretched.

Policy guardrail · Keep interventions proportionate to disruption

The effectiveness of any contrail mitigation policy should be assessed against both the climate benefit it delivers and the operational disruption it creates. Existing airline and airport systems have limited flexibility to change departure times at short notice, particularly where flight planning and coordination begin months in advance. Any intervention should therefore be proportionate to the expected benefit and supported by sufficient planning time, coordination and evidence of practical feasibility.

4.3.4 Assessment against the policy criteria

Measured against the Safe, Effective and Feasible criteria (see Section 2.2):

Assessment

- **Safe**: Scheduling avoidance changes when flights operate rather than how they are flown, so it introduces no new in-flight risk. Shifting flights into busier daytime periods increases pressure on runway and airspace capacity, which must be managed through existing ATC and airport processes.
- **Effective**: Static night-time restrictions rest on one of the better-established results in contrail science, but they are imprecise: most rescheduled flights would not have formed warming contrails, and daytime congestion may increase CO_2 emissions. Dynamic scheduling is more targeted, but the delays achievable within the take-off window are too short to move flights out of ISSRs, while longer delays cause substantial disruption and cost.
- **Feasible**: Static restrictions could build on existing night movement quotas, but require long lead times because they must be set at the strategic planning stage. Dynamic scheduling would require real-time monitoring and coordination between airlines, ATC and airports, and there is currently no evidence that it is feasible within existing operations.

Near-term role: complementary at most. Static night-time restrictions could reinforce route-based measures where night-time warming is material, but scheduling avoidance is not viable as a stand-alone instrument, and dynamic scheduling has no near-term role.

4.4 Fuel composition

Modifications to the aircraft exhaust composition through fuel alterations are a plausible lever for contrail mitigation because the particles emitted during combustion govern ice crystal formation in the exhaust plume. Fuels with lower aromatic content such as SAF, reduced-aromatic fossil kerosene and certain fuel additives produce fewer soot particles on combustion, reducing the number of available ice nucleation sites. This has led to proposals that wider SAF deployment could reduce contrail climate impact independently of flight routing [46].

This section assesses whether fuel composition is a viable alternative to other mitigation measures. Its conclusion is that SAF cannot currently be relied on as a meaningful contrail mitigation instrument under existing operational frameworks. The science is unresolved, supply is inadequate for fleet-scale deployment, and the commercial infrastructure through which SAF is traded is structurally incompatible with flight-specific contrail attribution.

This does not diminish the importance of understanding how SAF affects contrail formation and subsequent warming. SAF may still form part of a targeted measure (see Section 4.4.2). Although SAF should not be treated as the primary contrail mitigation measure, its growing use could affect contrail formation and forecast accuracy, so those effects need to be better understood.

4.4.1 The science is unresolved

SAF production routes often produce synthetic paraffinic kerosene (SPK). This has a lower aromatic and sulfur content than conventional jet fuel, which reduces soot particle emissions on combustion. On current rich-burn engines, fewer soot particles means fewer contrail ice crystals. In-flight experiments measured reductions of 50–70 per cent with blends [47] and 56 per cent with unblended HEFA-SPK [48]. However, the link between soot and ice crystal number is not straightforward: not all soot particles activate into ice, and volatile particles can contribute significantly [49]. Measurements behind a lean-burn engine found ice crystal numbers stayed high despite very low soot emissions [50], so the benefit of SAF may change as engine technology changes.

The effect on contrail radiative forcing is more uncertain than the effect on ice crystal numbers. With fewer crystals, each grows larger, so optical depth falls less than proportionally, although larger crystals also sediment faster and shorten contrail lifetime. Modelled reductions in contrail forcing for 100 per cent SAF range from about 26 per cent [48] to 44 per cent [21]. Separately, SAF's higher hydrogen content raises the Schmidt-Appleman threshold temperature, allowing contrails to form in slightly warmer conditions; this increases contrail occurrence by around 5 per cent for 100 per cent SAF in a model [21]. The resulting change in contrail radiative forcing cannot yet be generalised across blend ratios, SAF pathways, engine technologies or atmospheric conditions. Quantitative relationships exist in contrail models but have not been validated well enough for operational or policy use.

Policy guardrail · Research SAF effects regardless of mitigation pathway

As many countries implement SAF mandates and scale up production, it is important to understand how SAF affects contrail formation and subsequent warming estimates, regardless of the chosen mitigation pathway. As the use of SAF grows, its effects may become more significant. Many mitigation measures, including navigational avoidance, rely on contrail prediction models that may be affected by fuel properties. Additionally, if blended SAF is shown to significantly reduce contrail warming, fuel composition could be used as a complementary measure, especially when other measures are operationally infeasible due to congestion or safety concerns.

4.4.2 Supply constraints make fleet-scale impact implausible

Setting the science aside, the volumes required for SAF to materially shift fleet-average contrail forcing do not exist and are unlikely to become available within the current decade. Global SAF production currently represents well under 1 per cent of aviation fuel demand, dominated by hydroprocessed esters and fatty acids (HEFA), which face a hard feedstock ceiling due to competition with road transport biofuels and other sectors. The most ambitious national mandates, including those of the UK, target around 22 per cent SAF by volume by 2040 [51]. Alternative pathways such as power-to-liquid face cost and renewable electricity resource constraints that make meaningful scale contingent rather than assured well beyond that horizon. Even if the contrail benefit is real and unambiguous, the absolute effect on fleet-average contrail forcing at those blend ratios is modest at best.

Supply scarcity also determines how any future SAF contrail instrument would have to be scoped. Applied thinly across the whole fleet, SAF would achieve negligible per-flight impact; the same volume concentrated on the flights most likely to form warming contrails, such as those crossing ice-supersaturated regions at night or in high-impact seasons, could in principle achieve a meaningful localised reduction [46], [52]. This would require physical delivery of SAF to specific flights on specific days, which current supply chain infrastructure cannot provide. The general principle is that the thinner the supply of a mitigation input, the more essential targeting becomes, because thin universal application is indistinguishable from no policy at all.

Policy guardrail · Do not build contrail policy on SAF supply

SAF mandates are insufficient to produce material fleet-average contrail forcing reductions, even under optimistic efficacy assumptions. Contrail policy instruments must not be designed around SAF supply trajectories as a primary delivery mechanism; any framework that does so should be required to demonstrate quantified forcing impact at the blend ratios realistically achievable within the relevant compliance period.

4.4.3 The book-and-claim model is incompatible with contrail crediting and actively degrades contrail modelling

The third and most structurally fundamental problem is commercial. Under the book-and-claim model, the mechanism through which almost all SAF is currently traded, an airline purchases a certificate representing that SAF has been produced and supplied into the aviation fuel system somewhere, without that fuel being physically linked to the purchasing airline's own aircraft. This model works for CO_2 accounting because CO_2 is a well-mixed global gas, i.e. a tonne of emissions avoided anywhere is climatically equivalent to a tonne avoided elsewhere. Contrail forcing is not like this. It is hyper-local, depending on the specific exhaust composition of the specific flight in the specific atmospheric column at the specific moment. A flight that takes off with conventional fuel in its tanks produces conventional-fuel contrails regardless of what certificates the operating airline holds.

In practice, even where SAF is physically delivered into airport fuel systems, blending, storage and distribution infrastructure means the SAF fraction in any given aircraft's tanks at departure is unlikely to match the fraction the operator has claimed or purchased. An airline holding certificates for 50 per cent SAF on a given flight may have uplifted fuel containing 5 per cent, 20 per cent or 50 per cent SAF at that airport on that day, depending entirely on local supply. For CO_2 accounting this discrepancy is irrelevant. For contrail modelling it is not.

This matters beyond accounting, because credible frameworks for contrail mitigation, such as forecast-driven routing, ATM-level avoidance and market-based incentives for avoiding ice-supersaturated regions, depend on models that predict contrail formation and persistence for specific flights under specific atmospheric conditions. Those models require knowledge of the fuel actually burned, since exhaust composition governs the formation physics. The SAC threshold shift described above is a function of the water vapour emission index, which tracks the physical blend fraction, not the certificated one. A model using certificated fractions will systematically misestimate contrail formation thresholds. These errors will scale with the gap between certificated and physical blend fractions which will widen as book-and-claim volumes grow.

The result is an anti-synergy: the SAF certificate infrastructure, appropriate for its intended CO_2 accounting purpose, introduces a structural source of error into the contrail avoidance systems that

represent aviation's most practically viable path to non-CO_2 climate mitigation. Two instruments that are individually coherent become mutually degrading when deployed together.

Policy guardrail · Require physical fuel records, not certificates

Any contrail-related mitigation credit or compliance pathway that incorporates SAF blend assumptions must require verified physical fuel delivery records for the specific aircraft and departure, not certificate holdings. Regulators designing contrail incentive schemes must explicitly prohibit book-and-claim SAF certificates from serving as evidence of exhaust composition for contrail modelling or crediting purposes. Where physical blend fraction cannot be verified at the individual flight level, SAF shall not be treated as a variable in contrail formation or forcing calculations.

What about non-SAF methods?

SAF was the focus of this section due to its growing use in accordance with the EU and UK SAF mandates. However, blending SAF with fossil kerosene is not the only way to reduce the soot content in aircraft exhausts and thus mitigate contrail formation. The key is to minimise the aromatic content of the fuel - another way of doing this is through hydroprocessing.

Hydroprocessing treats fossil-fuel-based kerosene by introducing hydrogen under elevated pressure and temperature to remove impurities like sulphur and convert aromatic compounds into saturated alkanes. It is already present as a fuel finishing process at most modern refineries (albeit at far milder conditions than those that would be required for contrail mitigation), and has the potential to be employed at an industrial scale, potentially avoiding the supply constraints faced by SAFs. However, the literature on the contrail impact of hydroprocessed fuels is very sparse, and the book-and-claim model constraint would still apply to hydroprocessed fuel. In addition, the lifecycle emissions of this fuel could be significantly higher than its traditional form if the hydrogen used is not green [53].

4.4.4 Assessment against the policy criteria

Measured against the Safe, Effective and Feasible criteria (see Section 2.2):

Assessment

- → **Safe**: SAF is used as a drop-in fuel within certified blend limits, so it raises no new operational safety concerns.
- → **Effective**: The net effect of SAF on contrail radiative forcing is unresolved and may differ by blend ratio, production pathway and atmospheric conditions. Even if a benefit were established, current and mandated supply volumes would provide low coverage, and book-and-claim trading prevents the benefit from being attributed to individual flights. Targeted use could improve the effectiveness of scarce SAF if its efficacy and physical delivery can be demonstrated.
- → **Feasible**: Supply is constrained by feedstock availability and cost well beyond the near term. Physical blend fractions are not currently verified at the individual-flight level, as a contrail compliance pathway would require. Targeted use of available SAF may become operationally viable if supply can be physically delivered and tracked.

Near-term role: Complementary research and decarbonisation, not contrail crediting. SAF remains important for lifecycle CO_2 reduction, but contrail policy should not be designed around it or allow SAF contrail claims to substitute for routing-based obligations until the science, supply and attribution problems are resolved.

4.5 Aircraft design modifications

Advancements in engine technology and adaptations to existing aircraft designs present an alternative to purely operational contrail avoidance. This approach is generally termed *technical contrail suppression* and aims to prevent, or substantially reduce, contrail formation at the point of origin. Unlike navigational or scheduling strategies, technical contrail suppression avoids reliance on accurate meteorological prediction, airspace integration and monitoring, reporting and verification (MRV) attribution. In the most optimistic case, a fleet-wide adoption of technologies that inherently prevent contrail formation is the most effective method of contrail avoidance in the long term.

However, near-term feasibility is constrained by the low technology readiness level (TRL) of candidate engine modifications and by certification timescales intrinsic to commercial aviation. Novel engine technologies require sustained investment to advance TRL to the point of certified, market-ready options. Even at high TRL, the certification burden scales with modification scope: minor engine changes typically require only engine-level certification, whereas modifications interacting with airframe systems trigger dual certification of both engine and aircraft type certificate, a process typically taking 5-7 years [54], [55]. Beyond certification, fleet-wide penetration of new airframe-engine types is governed by aircraft

replacement cycles of 20-30 years [56]. The technical contrail suppression approach should therefore be treated as a parallel effort to operational strategies, not a substitute for them [57].

4.5.1 Classification and constraints

There are three main approaches to technical contrail suppression [27], [58] shown in Table 1:

Engine Design Approach	Effect
Reduce/remove water in exhaust gas	Limit amount of condensation
Reduce/remove particulate emissions	Limit condensation nuclei
Increase exhaust temperature	Limit the probability of reaching liquid-water saturation

Table 1: The different approaches to engine design for contrail avoidance and the methods they employ.

A full scientific review of each approach is beyond the scope of this primer. The following subsections summarise the current state of development within each method.

4.5.2 Water extraction and the water-enhanced turbofan

Studies indicate that, in standard kerosene engines, contrail prevention can be achieved at 100 per cent relative humidity if 76 per cent of water is removed from exhaust [59]. The most developed engineering proposal for reducing water vapour emissions at source was the water-enhanced turbofan (WET) concept [60], [61]. Initially planned for market launch by MTU Aero Engines in 2035, this project was discontinued in 2025 primarily due to its increased weight and fuel consumption [62], [63].

Condensation hardware must function reliably across the wide pressure, temperature and humidity envelope encountered during flight, and integration into a certified engine architecture represents a substantial engineering challenge. Near-term deployment is not credible under current development timelines.

4.5.3 Engine efficiency: a trade-off

A counter-intuitive but well-established result in contrail science is that improvements in propulsion efficiency, the central objective of modern engine development programmes, make contrail formation *more* likely [2]. This represents a direct and unresolved tension at the engine design level. Engine development programmes optimised for fuel efficiency and nitrogen oxides (NO_x) reduction do not naturally converge on configurations that minimise contrail formation propensity. Preliminary studies indicate that marginal efficiency reductions are not viable either. One design reduced the threshold formation temperature of contrails through variable guide vanes, but decreased overall efficiency by 10 per cent and increased specific fuel consumption by over 42 per cent [64]. Another approach, which studied increasing the bypass nozzle area, similarly found that specific fuel consumption increased by 36 per cent without achieving the goal of preventing persistent contrails [58].

4.5.4 Particulate reduction

The primary approach towards reducing particulate emissions is through fuel additives and fuel alternatives [65], previously covered in Section 4.4.

4.5.5 Alternative propulsion technologies

Longer-horizon propulsion technologies, principally liquid hydrogen combustion and electric or hybrid-electric architectures, have distinct implications for contrail formation that are not straightforwardly positive, and their development timelines place them outside the window relevant for near-term contrail policy.

Hydrogen (H_2) combustion eliminates CO_2 and soot emissions entirely, removing non-volatile particulate matter as a source of ice condensation nuclei, yet is expected to increase water vapour emission mass by 2.6 times compared to kerosene [66]. The true effect of hydrogen combustion on contrail formation remains an open research question with preliminary laboratory studies indicating that hydrogen contrail ice number and optical thickness may decrease, but form at lower altitudes and therefore increase contrail coverage [67].

Electric and hybrid-electric propulsion eliminates combustion exhaust from the propulsive system entirely, removing the thermodynamic basis for contrail formation at source. However, current and near-term battery energy density confines viable electric propulsion to short-haul operations at lower cruise altitudes, where flight paths are less likely to intersect with ice-supersaturated regions and contrail formation is correspondingly less probable [68]. For the large commercial aircraft responsible for the majority of contrail radiative forcing, electric propulsion has no credible near-to-medium-term pathway.

Both hydrogen and electric architectures require infrastructure transformation well beyond the aircraft itself, and fleet penetration at meaningful scale is not plausible before the second half of the century. Contrail policy should not be premised on their availability.

4.5.6 Assessment against the policy criteria

Measured against the Safe, Effective and Feasible criteria (see Section 2.2):

Assessment

- → **Safe**: Any engine or airframe modification must pass existing airworthiness certification before entering service, so safety is assured through established processes. The same processes are a major source of delay.
- → **Effective**: In the long term, technologies that prevent contrail formation at source could be the most effective form of mitigation, as they do not rely on forecasts or per-flight decisions. However, current approaches carry large fuel-burn penalties, improvements in propulsion efficiency make contrail formation more likely, and the net contrail effect of hydrogen propulsion remains uncertain.
- → **Feasible**: Candidate technologies are at low TRL; certification can take 5–7 years and fleet replacement cycles of 20–30 years mean meaningful fleet penetration is decades away. Aircraft design modifications are therefore not feasible as a near-term contrail policy instrument.

Near-term role: long-term research. Technical contrail suppression warrants investment as a parallel effort, but contrail policy designed around it would defer action on a problem with near-term climate relevance.

Policy guardrail · Treat aircraft design as long-term

Aircraft modifications should be treated as a longer-term research and investment pathway, not as the primary instrument for near-term contrail policy. Certification and fleet-renewal timescales mean that meaningful deployment will take decades; operational measures should therefore address near-term climate risk while technical solutions continue to mature.

4.6 Comparing mitigation mechanisms

Table 2 records when each mechanism acts in the flight-planning process, who makes the decision and how far ahead of the flight that decision is made (Section 4.1.1).

Mechanism	Phase	Decision-maker	Lead time
Pre-tactical navigational	Pre-tactical	Airline dispatchers	12–24 hours
Tactical navigational	Tactical	ATC and crew	Minutes to real time
Static scheduling	Strategic	Network planners and slot coordinators	Months
Dynamic scheduling	ATC clearance	ATC	Minutes
Fuel composition (SAF)	Strategic supply; pre-tactical uplift	Fuel suppliers and airlines	Years (supply); hours (uplift)
Aircraft design	Strategic	Manufacturers and regulators	Years to decades

Table 2: Operational mechanisms for contrail mitigation by intervention phase, decision-maker and lead time.

Table 3 then qualitatively summarises how each mechanism performs against the Safe, Effective and Feasible criteria (see Section 2.2), drawing on the assessments in the preceding sections.

Mechanism	Safe	Effective	Feasible	Near-term role
Pre-tactical navigational	Yes: planned deviations cleared through ATC	High: targets the small share of high-warming flights; small fuel penalty	High: existing aircraft and ATM; tooling needs wider deployment	Primary
Tactical navigational	Yes, but adds controller and crew workload	High: responds to actual conditions; limited by capacity	Medium: workload and airspace capacity limit scale	Complementary
Static scheduling	Yes: no change to how flights are flown	Low: imprecise targeting; daytime congestion may add CO_2	Medium: builds on night quotas, but needs long lead times	Complementary at most
Dynamic scheduling	Yes: no change to how flights are flown	Low: feasible delays too short to avoid ISSRs	Low: no evidence of operational feasibility	None
Fuel composition (SAF)	Yes: certified drop-in fuel	Unproven: net forcing effect	Low at scale: constrained	Complementary research

Mechanism	Safe	Effective	Feasible	Near-term role
		unresolved; low coverage	supply; blend not verifiable per flight. Medium for targeting	
Aircraft design	Yes, via airworthiness certification	Potentially high long-term; current designs carry large fuel penalties	Low: low TRL, long certification and fleet-renewal timescales	Long-term research

Table 3: Assessment of operational mechanisms for contrail mitigation against the Safe, Effective and Feasible policy criteria.

The comparison also shows why timing matters (Figure 3). Mechanisms fixed at the strategic stage (static scheduling, aircraft design) are set before any useful atmospheric forecast exists, so they cannot target the specific flights that cause most warming. Tactical avoidance responds to actual conditions but transfers the decision to ATC at the point where capacity is most constrained (Section 5.5). Pre-tactical navigational avoidance is the only mechanism that acts when a usable forecast first exists, through an existing decision-maker (the airline dispatcher), using an existing artefact (the filed flight plan).

This is why pre-tactical navigational avoidance is the only mechanism rated high on both effectiveness and feasibility today, making it the primary candidate for near-term policy. Scheduling avoidance can complement it in limited cases, while SAF and aircraft design modifications sit outside the near-term policy window. Section 5 tests what limits navigational avoidance in practice.

5 Technical constraints on navigational contrail avoidance

Section 4 identified pre-tactical navigational avoidance as the primary near-term mechanism. This section examines what limits it in practice, and which of those limits policy can remove.

5.1 Constraint taxonomy

Not all constraints on contrail avoidance are the same kind of problem. Some can be overcome through deliberate investment, coordination or policy. Others reflect structural features of the atmosphere, of air traffic management or of international governance that no amount of funding or patience will eliminate. Policy designed without this distinction risks either premature mandates (demanding precision that physics cannot deliver) or indefinite delay (waiting for breakthroughs where targeted investment is the actual bottleneck). This report classifies constraints into the two categories shown in Table 4:

Category	Character	Policy Response
Investment-dependent	Can be improved through deliberate funding and coordination—for example, by investing in humidity-observation networks, model resolution, contrail-aware flight-planning tools, MRV infrastructure and interoperability standards.	Identify specific interventions; commit resources; track progress; avoid locking in today's limitations.
Structurally persistent	Reflects fundamental limits—for example, chaotic atmospheric prediction at small scales, the multi-actor attribution problem and international governance architecture.	Design policy around these constraints rather than waiting for them to resolve.

Table 4: The two categories of contrail avoidance constraints, their characteristics and the appropriate policy responses.

5.1.1 System dependencies

Implementing Safe, Effective and Feasible contrail management policy is further challenged by system dependencies. Observational infrastructure, meteorological data, decision-support systems, operational procedures and regulatory efforts all depend on and inform one another. Figure 6 illustrates this stack.

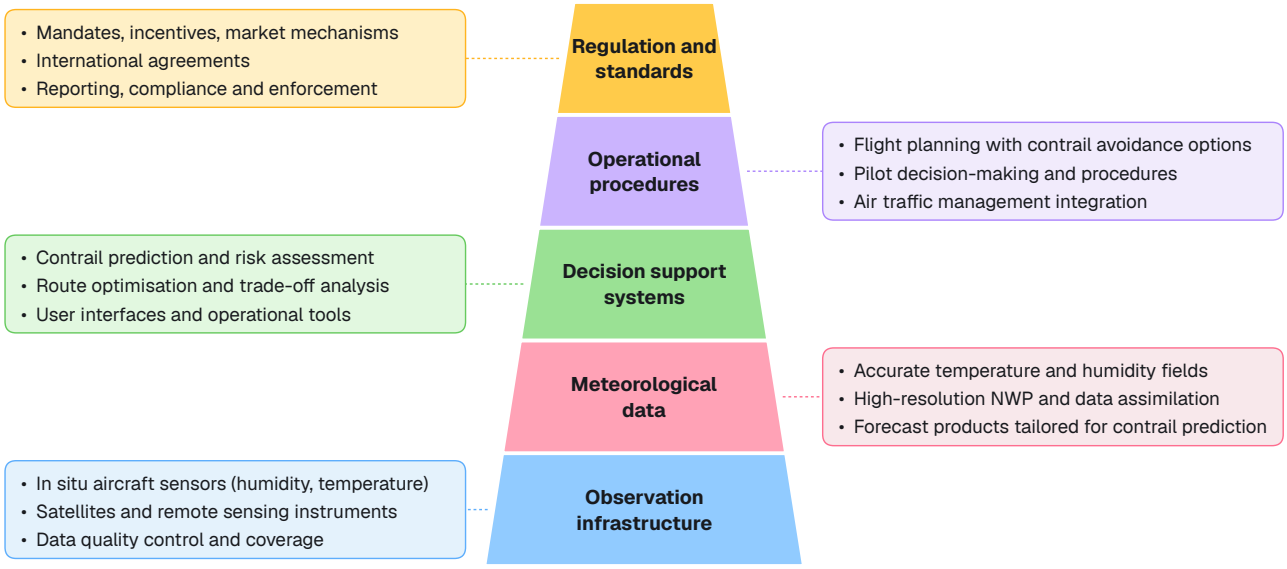


Figure 6: Diagram of the stack of system dependencies for contrail avoidance.

The stack features bidirectional interactions: each layer depends on those below it and informs those above it. Weakness at any layer constrains the layers above it, so policy that targets only the regulatory layer without ensuring the lower layers are in place will produce mandates that cannot be met. An in-depth understanding of the foundation layers is essential to designing procedures that work within structurally persistent constraints and target investment in areas essential to policy development. The following sections describe each layer and the constraints associated with it.

5.2 Observation infrastructure

This layer supplies the measurements of temperature, humidity and contrails on which every layer above it relies, making it the foundation of the stack. Coverage can be expanded through investment, but limits on how many contrails passive satellites can detect, and how soon after formation, are structurally persistent. These limits carry through to attribution, discussed under Section 5.6.1.

5.2.1 Humidity and temperature sensing

Investment-dependent

Temperature observations at cruise altitude are generally adequate, but upper-tropospheric humidity observations (a critical variable for contrail persistence) remain sparse [25]. Much of the humidity information used by forecast models is derived from indirect estimates rather than direct measurements. The gap is addressable: enhanced humidity sensing (aircraft-based water-vapour sensors, improved radiosonde coverage and next-generation satellite instruments) has a known solution pathway, and programmes such as In-service Aircraft for a Global Observing System (IAGOS) and emerging commercial aircraft sensor networks demonstrate its feasibility [69]. How these observations feed into forecasts is covered in the next layer (Section 5.3).

Policy guardrail · Prioritise investment in humidity observation

Observation gaps in humidity sensing are among the most addressable bottlenecks for contrail avoidance and should be treated as an explicit investment priority. Appropriate measures could include mandating or incentivising aircraft-based humidity sensors on commercial fleets and funding dedicated upper-tropospheric observation campaigns.

5.2.2 Satellite detection of contrails

Investment-dependent | Structurally persistent

Direct observation of contrails is possible using satellite imagery, ground-based lidar and aircraft-mounted sensors [70], but global coverage by these sensors is incomplete. Current geostationary satellites, the most common method for observing contrails, can detect only approximately 46 per cent of contrails at any given moment. This observation ceiling is primarily imposed by the optical thickness of the remaining contrails rather than by the resolution of the satellite images themselves [71]. Consequently, the ceiling is largely insensitive to resolution improvements: increasing the pixel resolution from 2 km, typical of current-generation satellites, to 0.5 km raises the detectable fraction by approximately 4 percentage points, leaving about half of contrails undetectable by satellites. Investment in next-generation satellite constellations, machine-learning detection algorithms and ground-based sensor networks will improve observational capability at the margins, but no identified investment pathway closes the fundamental optical-thickness gap.

Figure 7 shows the fraction of contrails observable by satellite as a function of age. Observability is lower immediately after formation, when contrails are narrow or optically thin, and peaks at around one hour before declining as contrails dissipate. Increasing satellite resolution from 2 km to 0.5 km improves the detection of young contrails, but has almost no impact on chance of detection after 2 hours.

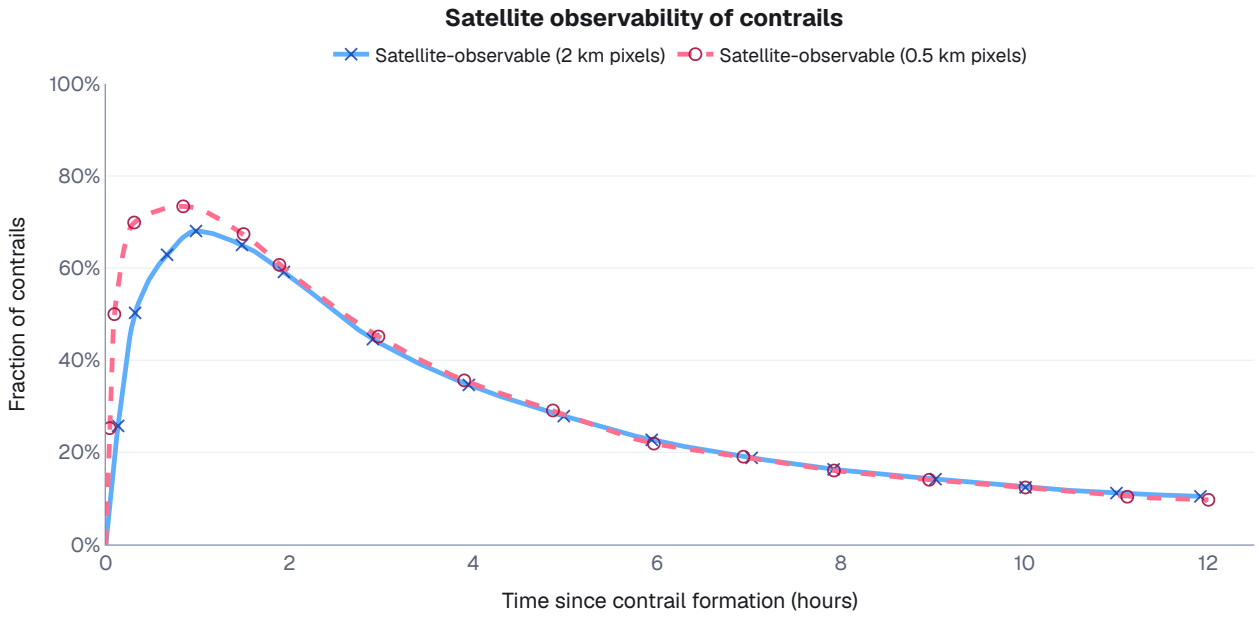


Figure 7: Fraction of contrails observable in satellite imagery at 2 km and 0.5 km pixel resolution by time since formation, using data from Driver et al. [71]. The higher resolution primarily improves observability for younger contrails.

The observation estimates above assume clear-sky conditions: in practice, overlying clouds can hide contrails from satellites entirely. Other contrails formed in the same region can obstruct the linear shape that observation algorithms usually look for, preventing accurate observation.

Geostationary satellites first detect contrails 10–45 minutes after formation [71], [72], [73], with a mean first detection time of 41 minutes [74]. This is not a sensor quality problem. It is a consequence of the physics of contrail formation: contrails must spread laterally before becoming detectable.

This delay in contrail detection has direct implications for which policy measures are feasible. The detectable fraction is not a random sample of all contrails: optically thick, longer-lived contrails tend to be both more detectable and to contribute disproportionately to total radiative forcing. Therefore, the 46 per cent of contrails observable at any given moment may represent around 70 per cent of total energy forcing (EF) [71].

Ground cameras detect contrails much faster than satellites but are geographically concentrated over land in Europe, North America and Australia, leaving oceanic routes entirely uncovered. A purpose-built low-Earth-orbit constellation with short revisit times (~5 minutes) would substantially close the detection-lag gap, but no such system is currently planned for contrail detection, and the optical ceiling would still apply to any passive infrared instrument.

To conclude, 46 per cent of contrails is the estimated observational ceiling at a given instant in the atmosphere with current technology; it does not account for cloud occlusion or other sources of uncertainty. However, the contrails that can be detected are larger contributors to total radiative forcing. Even though not every contrail can be observed, observational data can still be used to validate methodologies and improve contrail-formation models. Section 5.6 explores the implications of these observation constraints for MRV.

Policy guardrail · Do not expect better satellites to see every contrail

Policy should not assume that better satellites will make individual contrails observable: the optical-thickness ceiling means roughly half of contrails will remain undetectable by passive satellite instruments regardless of investment.

5.2.3 Role in the dependency stack

→ **Depends on**: funding and sensor deployment only.

→ **Informs**: meteorological data (layer 2) through humidity and temperature observations, and regulation (layer 5) through the observational datasets used to validate models.

5.3 Meteorological data

This layer turns raw observations into the temperature and humidity fields, forecasts and tailored products on which contrail prediction depends.

5.3.1 Model resolution and the scale of ice-supersaturated regions

Investment-dependent | Structurally persistent

Effective contrail avoidance requires accurate information about atmospheric conditions at cruise altitude, particularly temperature and humidity. The ISSRs where persistent contrails form are typically described as pancake-shaped: tens of kilometres wide, vertically shallow (hundreds of metres), highly dynamic and difficult to observe directly [75].

Studies have identified coarse model resolution, unresolved moisture variability and limitations in the representation of ice supersaturation as primary sources of forecast error in contrail prediction [25].

Existing meteorological systems were designed for weather forecasting, not contrail mitigation, and this limitation is compounded by a resolution mismatch. Even world-leading weather forecast models, such as the European Centre for Medium-Range Weather Forecasts (ECMWF) Integrated Forecasting System (IFS), run at approximately 9 km horizontal resolution, with vertical level spacing of roughly 300–400 m at cruise altitude. Widely used reanalysis products for this forecast, such as the ECMWF Reanalysis version 5 (ERA5), are coarser still, at approximately 31 km horizontally [76]. Contrail formation depends on moisture structures at substantially smaller scales: ISSRs just a few kilometres wide and hundreds of metres deep may span only a handful of grid cells, while sub-grid humidity variability remains unresolved. Figure 8 illustrates this resolution mismatch: the median ISSR is approximately 3 km long [77] and 165 m deep [78] —significantly smaller than ERA5 or ECMWF IFS grid cells—meaning that many persistent contrail-forming ISSRs are missed.

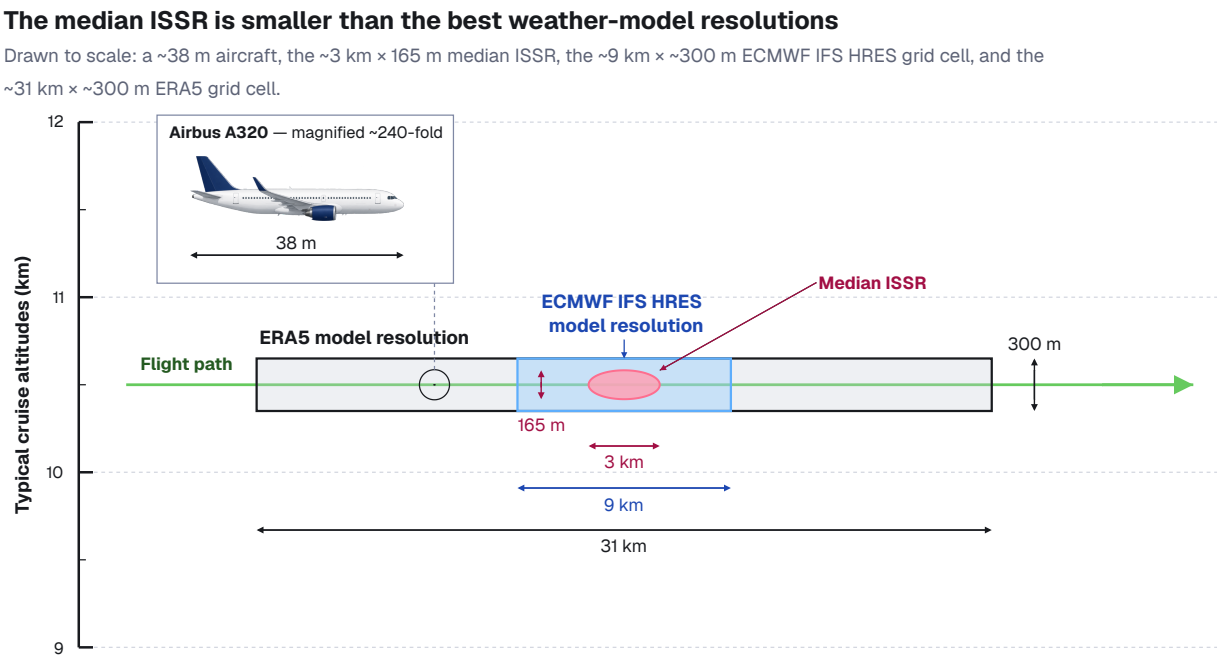


Figure 8: Illustration of the relative sizes of an A320, the median contrail-forming ISSR and the resolution of the weather model grid cells commonly used for forecasting contrails. Note that the vertical scale is ten times larger than the horizontal scale — a true-to-scale depiction would make the ISSR and grid cells appear as extremely thin rectangles.

The purpose of the illustration above is to visualise the complexity of the problem, not to suggest that identifying ISSRs is operationally infeasible. As discussed in the next section, Section 5.4.1, a hit rate of 60 per cent is achievable with current model resolution and available research. Although predicting every ISSR may be difficult at this resolution, that does not rule out the operational feasibility of predicting most warming contrails. The most warming contrails are likely to form in ISSRs larger than the median: these larger regions give contrails more opportunity to grow and persist, increasing their warming impact.

This is a capability gap with a known solution pathway: higher-resolution numerical weather prediction and dedicated data assimilation for the upper troposphere, supported by the enhanced humidity sensing described in Section 5.2.1. Realising operational benefits would require years of sustained investment in observation infrastructure and model development before forecast skill could reach the level needed for reliable flight-level contrail prediction. Even with higher resolution, some humidity variability will remain below what observations and models can resolve. This part of the constraint is thus structurally persistent.

5.3.2 Role in the dependency stack

→ **Depends on:** observation infrastructure (layer 1), particularly humidity measurements.

→ **Informs:** decision support systems (layer 3), whose prediction skill is bounded by the quality of these fields.

Policy guardrail · Invest in contrail-relevant weather forecasting

Meteorological forecast quality is an explicit investment priority alongside observation. Appropriate measures could include funding higher-resolution and contrail-tailored NWP products, supporting data assimilation for the upper troposphere and establishing forecast-skill benchmarks to inform future policy activation. A framework that recognises both current limitations and the need for sustained improvement can support progress without presuming that better capability will emerge on its own.

5.4 Decision support systems

This layer uses meteorological data to predict contrail formation, persistence and climate impact and to generate route recommendations that account for fuel burn and cost.

5.4.1 Contrail prediction and risk assessment

Investment-dependent | Structurally persistent

Contrail avoidance depends on predicting where contrails will form, how long they will persist and their resulting climate impact. These predictions combine weather forecasts with aircraft-performance models and contrail-evolution algorithms. Scientific understanding of contrail formation (see Section 3.1) has advanced considerably, but two distinct types of uncertainty remain, and they demand different policy responses.

Model error (investment-dependent): Current models make approximations in their treatment of ice nucleation, contrail spreading, optical properties and radiative transfer [79]. These approximations introduce systematic biases that differ between modelling frameworks. Two models assessing the same flight may produce materially different estimates of contrail radiative forcing depending on their assumptions and input data. This type of error can be reduced through model intercomparison, observational validation and algorithmic refinement. Model agreement could narrow over time, provided investment in validation infrastructure continues.

Fundamental predictability limits (structurally persistent): Even with perfect models, contrail prediction at the individual-flight level faces irreducible uncertainty from atmospheric chaos. Small-scale humidity variability that determines whether a specific flight forms a persistent contrail may never be predictable beyond probabilistic bounds. The atmosphere at cruise altitude exhibits sensitive dependence on initial conditions at scales below observational resolution. No foreseeable improvement in modelling or observation is likely to eliminate this floor of uncertainty for an individual flight.

Aggregate prediction can reduce this uncertainty at the portfolio level. A model that cannot reliably determine whether flight BA283 on Tuesday will form a warming contrail can still estimate, with reasonable accuracy, what fraction of flights through a given region and altitude band will produce net-

warming contrails over a season. This is the same distinction that applies in climate modelling, where individual weather events cannot be predicted beyond a few days but seasonal and regional statistics can be estimated with useful skill.

Contrail models are therefore decision-support tools rather than precise measurement systems. Their value lies in identifying likely mitigation opportunities at portfolio scale, not in certifying the climate impact of individual flights.

Figure 9 is from ContrailBench [80], a new framework for evaluating contrail forecasts against real-world observations. The hit rate on the y-axis is the proportion of true contrail regions correctly forecast by Contrails.org's model, while the x-axis is the fraction of flight distance forecast to form persistent contrail regions (PCRs), and therefore subject to avoidance measures. This fraction can be viewed as a proxy for cost: the greater the share of flight distance forecast to produce contrails, the greater the fuel burn and flight disruption associated with avoidance. The figure shows two curves, representing different observation sources: weather balloon measurements from the Global Climate Observing System (GCOS) Reference Upper-Air Network (GRUAN) and IAGOS in-flight measurements. The data points are produced by varying the horizontal buffer (or prediction margin) added to contrail forecast regions. The leftmost data point is the raw forecast, without any buffer, and achieves a 60 per cent hit rate. Each subsequent data point expands the forecast region; a larger buffer means a higher hit rate, but the trade-off is that avoidance disrupts more flights. The rightmost data point uses a buffer that expands the raw forecast area by a factor of 100. This achieves a 93 per cent hit rate but requires 24 per cent of global flight distance to be rerouted. The IAGOS PCR rate indicates that measurements show contrails forming over 10 per cent of flight distance. The closest model forecast data point matches this share with a small buffer, achieving a 70 per cent hit rate. By contrast, randomly guessing which regions of the atmosphere form contrails would achieve a 10 per cent hit rate. Thus, despite substantial scientific uncertainty, the models show clear skill over random guessing at the aggregate scale.

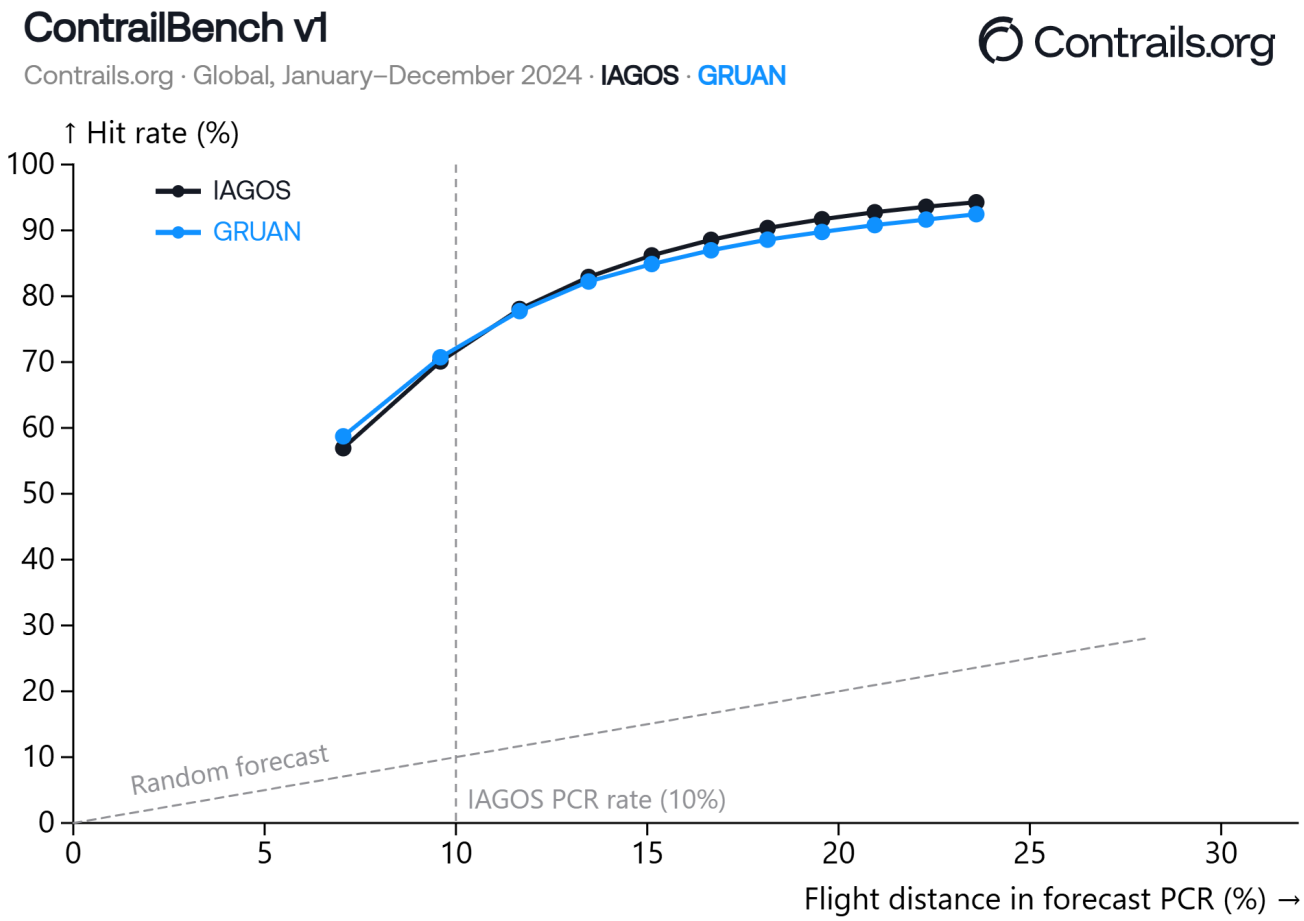


Figure 9: Graph from ContrailBench v1 [80] showing the relationship between the percentage of flight distance that is contrail-forming and the hit rate of Contrails.org's forecast model compared with in-flight (IAGOS) and weather balloon (GRUAN) measurements. Adding larger safety buffers to predicted contrail-forming regions increases the hit rate, but at the cost of rerouting more flight distance.

5.4.2 Metric choice in trade-off analysis

Structurally persistent

The global warming potential (GWP) over a 100-year horizon (GWP100) is commonly used to quantify contrail warming and has been used in trials [38] and to estimate the contrail-related warming of individual flights [21], [28]. Different metrics can produce CO_2-equivalent values for the same contrail that vary by up to an order of magnitude. However, the variation is much narrower when the metric is applied to operational avoidance. Borella et al. [81] demonstrate this by making a case for contrail avoidance under nine different metrics. The size of the perceived benefit depends strongly on the metric, but the metrics agree on most avoidance decisions.

Figure 10 illustrates this for three absolute global warming potential (AGWP) time horizons. Moving from a 100-year to a 20-year horizon more than triples the estimated warming from contrails relative to CO_2, yet the share of contrail-forming flights for which avoidance is beneficial changes by only around two percentage points.

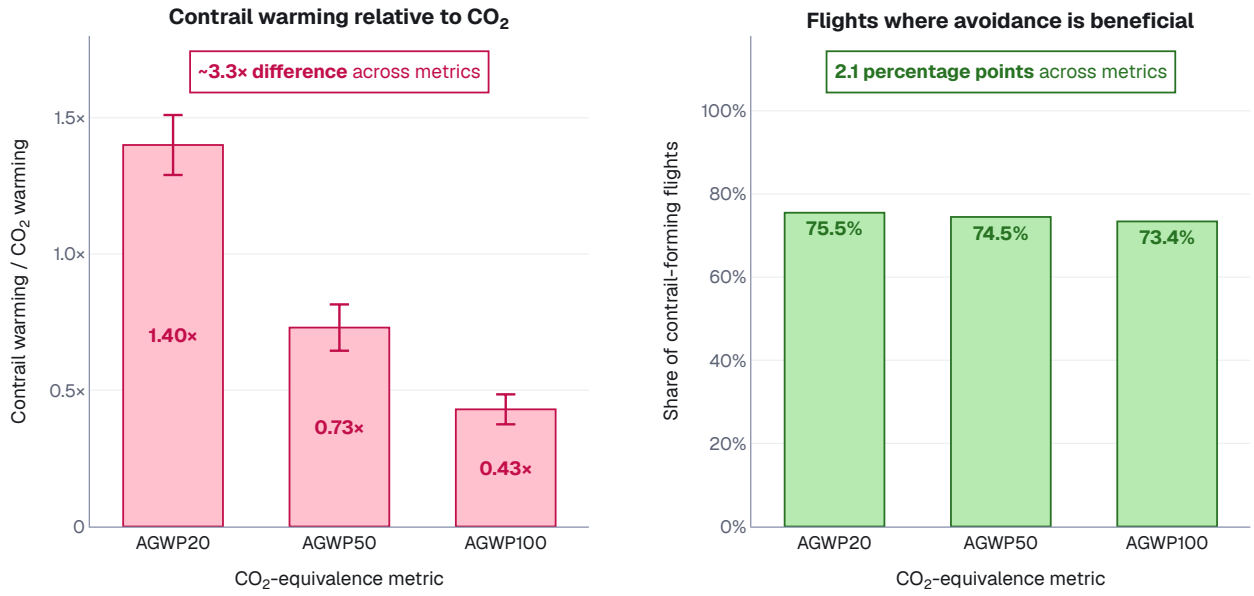


Figure 10: Effect of the CO_2-equivalence metric on contrail avoidance for North Atlantic flights in 2019, using data from Borella et al. [81]. Left: cumulative climate impact of contrails relative to that of the CO_2 emitted, under AGWP with 20-, 50- and 100-year time horizons. Right: share of the 260,854 contrail-forming flights for which rerouting is beneficial to the climate, assuming contrail formation is fully avoided at a 1 per cent fuel-burn (CO_2) penalty.

Because contrail warming is highly skewed across flights, with very few flights causing the majority of the warming (see Section 3.1.3), the flights worth avoiding are largely the same regardless of the metric applied. Metric choice does not decide whether contrail avoidance is beneficial overall. It decides which edge cases cross the threshold of perceived cost-benefit.

The effect of metric choice is best understood by considering the two extremes. If contrails caused no warming, no flight would be worth deviating, since any extra fuel burn would be a pure CO_2 penalty. If contrails were extremely warming, every flight forming a contrail would be worth deviating, whatever the fuel cost. The true cost-benefit lies somewhere between these extremes, and the metric determines where in that range the decision threshold sits. A metric that weighs contrail warming more heavily lowers the threshold, so more marginal flights are deviated and more fuel is burned. A metric that weighs it less raises the threshold, leaving some avoidable warming unaddressed.

The risk of choosing the wrong metric therefore lies at the margin rather than in the core case for avoidance. The flights with the largest contrail warming are deviated under almost any reasonable metric, but a metric that overstates contrail warming could lead to fuel being burned on edge cases where the extra CO_2 outweighs the benefit. If the science later revises contrail forcing, fuel penalties justified under the earlier metric may no longer be warranted.

Better science alone is unlikely to resolve this, because a perfect metric would require stakeholders to agree on how to prioritise short-term contrail warming against long-term CO_2 warming. The right metric is as much a political and economic choice as a scientific one.

5.4.3 Route optimisation

Investment-dependent

Flight-planning systems today optimise for fuel consumption, flight time and cost. Incorporating contrail climate impact requires additional modelling layers, real-time meteorological data feeds and multi-objective optimisation algorithms that trade fuel burn against contrail warming. These systems exist but are not widely deployed.

General improvements in flight-planning software will reduce the cost of adding contrail-aware optimisation, but will not by themselves produce interoperability standards or support for smaller operators.

Policy guardrail · Do not conflate CO_2 and non-CO_2 effects

There is ongoing debate about the most accurate way to compare the effects of CO_2 and contrails. From a policy perspective, the safest response to this uncertainty is to keep CO_2 mitigation and contrail mitigation as separate as possible, so that neither depends on an exchange rate that remains contested. In practice, this means contrail obligations should not be met with CO_2 reductions or allowances, nor CO_2 obligations with contrail avoidance, which would allow action on one to offset inaction on the other.

Avoiding this comparison in flight operations is harder. Flight-planning systems already trade off competing objectives, and contrail avoidance extends this by trading contrail warming against the CO_2 produced by additional fuel burn. A hard limit on where aircraft may fly for climate benefit lacks the flexibility the system needs to avoid unnecessary fuel burn and maintain safe separation. If, on a given day, up to 30 per cent of the airspace is covered by ISSRs, successfully deviating every aircraft may not be operationally feasible. Therefore, flight-planning software needs a way to optimise the overall impact of a flight across several parameters, including contrail warming.

5.4.4 Tools and capability asymmetry

Investment-dependent

Large-scale contrail avoidance requires technical capabilities that do not yet exist uniformly across the aviation ecosystem. Airlines, ANSPs, meteorological agencies and regulators each require new tools, processes and expertise. The constraint is not that these capabilities are impossible; rather, they require deliberate, coordinated development.

Capability maturity varies substantially across the ecosystem. Major carriers and ANSPs could likely develop bespoke optimisation systems; however, smaller operators will likely depend on externally provided services or standardised tools that do not yet exist. This asymmetry creates a risk that early

policy mandates produce compliance only among those already capable, while systematically excluding operators whose participation matters for aggregate climate outcomes.

Interoperability compounds the challenge. Contrail-avoidance decisions depend on shared assumptions about weather forecasts, climate metrics, model methodologies and reporting formats. Without common standards, different organisations assessing the same flight may reach different conclusions, undermining both market mechanisms and regulatory consistency. Standardisation bodies such as the European Organisation for Civil Aviation Equipment (EUROCAE) and the Society of Automotive Engineers (SAE) have not yet produced the required technical standards.

5.4.5 Role in the dependency stack

→ **Depends on:** meteorological data (layer 2) for forecasts and model skill in Section 5.4.1.

→ **Informs:** operational procedures (layer 4), which can use only the options that decision-support systems present.

Policy guardrail · Phase obligations in line with system readiness

Because capability constraints are investment-dependent rather than structurally permanent, a phased implementation tied to measurable readiness milestones is likely to be the most proportionate approach. Priorities could include funded development of open interoperability standards for contrail-aware flight planning, technical assistance for less-resourced operators and states and a capability-maturity framework that allows obligations to increase as infrastructure develops. The timing of any mandatory requirements should follow demonstrated system readiness rather than an arbitrary timetable.

Policy guardrail · Use aggregate, probabilistic performance metrics

Because individual-flight contrail prediction faces irreducible uncertainty from atmospheric chaos, policy should use aggregate and probabilistic performance metrics from the outset. This is best understood as a design requirement rather than a temporary compromise while science improves. Fleet-level or route-level performance targets, probabilistic scoring (such as expected contrail energy forcing reduction across a portfolio of flights) and risk-weighted model outputs would be more defensible than deterministic claims about individual flights. Model improvements may narrow uncertainty bands over time, but the underlying aggregate approach should remain central to the policy architecture.

5.5 Operational procedures

This layer covers how flights are actually planned, flown and managed once decision support is available: flight planning, pilot and controller procedures and integration with ATM.

5.5.1 Flight planning

Investment-dependent | Structurally persistent

Navigational avoidance is most readily implemented at the pre-tactical planning stage (see Section 4.2.1). Airline planning systems build a flight plan typically 12–24 hours before departure, dispatchers finalise the plan and then file it with the Network Manager or ANSP a few hours before departure (Section 8.1.2). This plan is optimised to minimise fuel burn, time and cost; contrail avoidance adds the additional objective of minimising contrail warming, informed by the forecasts and tools described in Section 5.4. The North Atlantic adds another layer of complexity to this process, where planning is tied to the Organised Track System (OTS), a dynamic network of unidirectional routes published twice daily by the Gander and Shanwick control centres. These tracks serve several purposes, including (i) maintain safe separation between aircraft where radar coverage is sparse; (ii) account for changing weather patterns to minimise fuel burn and flight time; and (iii) accommodate passenger traffic patterns and airline preferences.

Each step along the way of this process introduces the risk of abandoning navigational contrail avoidance. In the largest airline-led trial to date with American Airlines, dispatchers were presented with the choice of a contrail-optimised plan (for contrail-forming flights) and non-avoidance alternative plans. Dispatchers only released contrail-optimised plans for 15.4 per cent of eligible flights, and only 7.8 per cent were flown as planned [38]. The authors note that the voluntary participation and design of the trial were the primary causes of the low uptake, so these values indicate gaps in their operational setting rather than fixed ceilings. A common reason given for non-compliance was that, even though contrail avoidance plans sometimes involved descents and ascents in the middle of a flight that are safe and done in coordination with air traffic control, dispatchers and flight crews preferred to avoid these deviations. Tooling was part of this issue - OTS adherence features could not be used with contrail plans due to technical issues, dispatchers lacked access to vertical profiles of the atmosphere to visualise avoidance manoeuvres and pilots lack access to contrail information in their tools. These limitations can be overcome by investments into both better tooling and further trials to expand experience and familiarity with navigational avoidance.

However, the gap between planned and flown trajectories is partly structural. A filed flight plan is only a request: en-route clearances, weather and forecast updates can change between planning and departure, and these may not always propagate to the flown trajectory. Avoidance should be measured on the trajectories flown, not planned, but it would be unfair to penalise operators for failing to follow a contrail avoidance plan due to other constraints despite their intent to mitigate warming.

5.5.2 Variability across flights

Investment-dependent | Structurally persistent

Contrail avoidance requires changes to aircraft trajectories: altitude adjustments (typically ±2,000–4,000 ft, confirmed by operational trials) or lateral route modifications. The burden of avoidance can vary enormously between flights, depending on the airspace, congestion and size of the region being avoided. A flight in uncongested oceanic airspace may avoid a contrail-forming region with a 1,000 ft altitude change and negligible fuel penalty. A flight in dense European airspace may face a choice between a 4,000 ft deviation, a lateral reroute (which could increase fuel burn and CO_2 emissions by up to 5 per cent) or no avoidance at all. Although fleet-wide fuel burn penalty is modelled to be small (around 0.1–0.2 per cent of total fuel [30]), this will be concentrated in a fraction of flights, and some individual deviations may be costly.

Policy that treats all flights as equally capable of avoidance risks producing perverse outcomes. A uniform per-flight obligation would only be met where avoidance is cheap and easy, and would either exempt the hard cases, which may include flights with large warming, or impose large CO_2 penalties on them. The variability across flights is permanent, but its range can be narrowed as navigational avoidance matures and operational procedures evolve with better tooling and standards.

5.5.3 Dispatcher, crew and controller workload

Investment-dependent | Structurally persistent

Once a flight departs, any change to the filed plan must be agreed between the crew and a sequence of sector controllers (Section 8.1.3). Deviations for weather, separation and fuel are already routine, and turbulence avoidance shows that hazard-driven rerouting can be integrated into daily workflows (Section 2.1.3). Regardless of whether a contrail deviation is ANSP-led or airline-led (see Section 4.2.2), both parties must agree before it can take place. Workload depends on the number of aircraft requiring deviations and the level of traffic in the region.

Tactical avoidance (real-time rerouting during flight) is particularly constrained. Unlike pre-planned route optimisation, tactical deviations require immediate coordination between pilots and controllers. Each trajectory change increases controller workload and may trigger a cascade of adjustments across multiple aircraft to maintain separation. Real-time simulations at EUROCONTROL's Maastricht Upper Area Control Centre (MUAC) investigating the impact of contrail avoidance on capacity found that controllers prefer planned (pre-tactical) over tactical avoidance. Blocking flight-level bands for avoidance reduced capacity at moderate traffic and became impossible at high traffic [82]. At low traffic, such as at night, the capacity impact was low. Tactical avoidance is therefore possible, but its availability falls as traffic rises. This limit is structural: controller workload and separation requirements bound what can be done in real time, however good the tools become.

Pre-tactical avoidance shifts the decision away from controllers and aircraft crew to dispatchers and planning software, reducing overall workload by moving it upstream and increasing automation. In the

American Airlines trial, contrail-optimised plan uptake was reduced when workload was higher, and avoidance plans were set aside when turbulence was a concern. Their stated priorities were safety, then efficiency, then contrails [38]. These obstacles are investment-dependent; automated plan selection, integrated interfaces and common contrail information sources shared between pilots and controllers can ease the burden of avoidance.

5.5.4 Airspace capacity and rerouting envelope

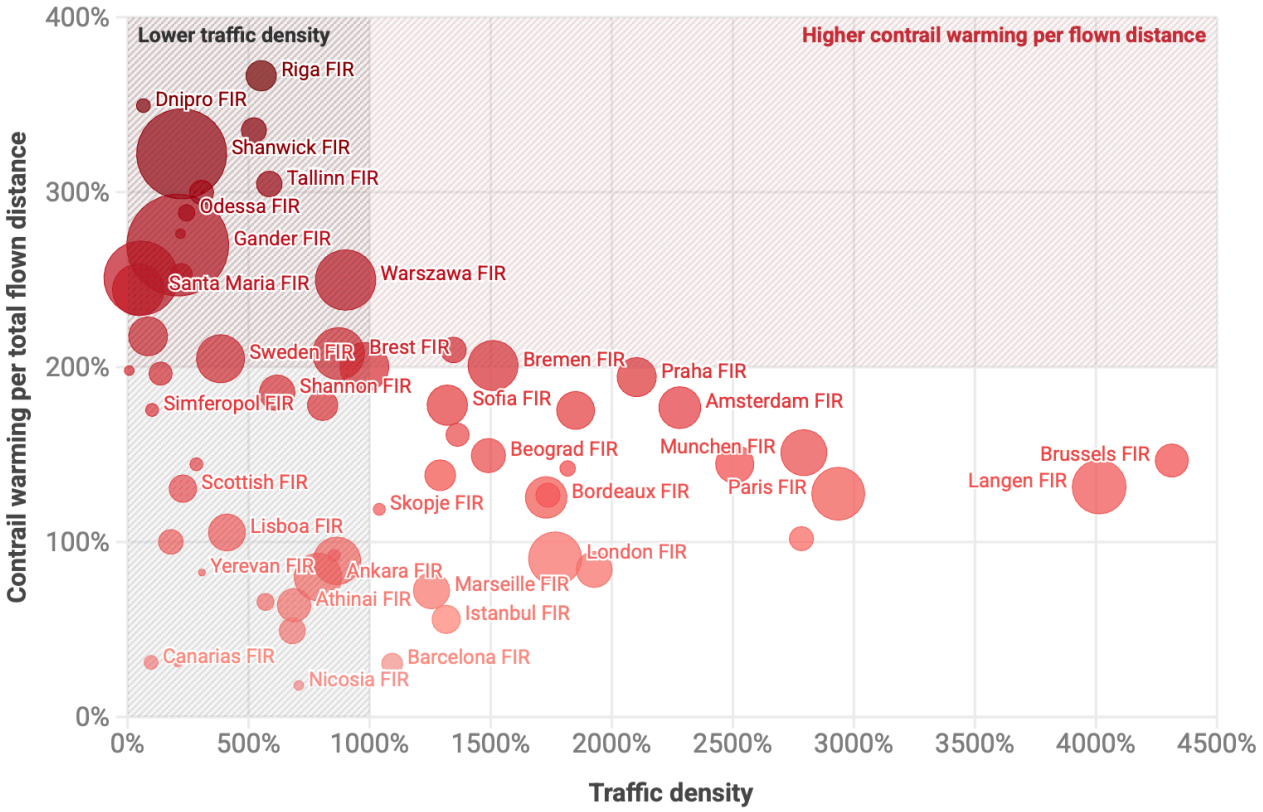
Investment-dependent | Structurally persistent

Any avoidance-related trajectory changes must be accommodated within ATM systems whose primary design objectives are safety, separation assurance and capacity. These objectives take precedence over climate optimisation, and contrail-driven routing competes with them for the same finite airspace. In congested sectors such as the North Atlantic OTS, European core area and US northeastern corridor, opportunities for vertical or lateral deviation are frequently constrained by traffic density, airspace design, military restrictions and weather avoidance routing. As an example, MUAC's simulations indicate capacity reductions of around 20 per cent in moderate traffic when avoidance is applied [82]. The practical rerouting envelope available on any given flight is often far narrower than modelling studies assume.

The figures below illustrate this constraint. Figure 11 plots contrail warming impact against traffic density (as percentages of the 2019 global averages) for European and North Atlantic airspaces. The airspaces with the greatest warming impact, such as the Shanwick, Gander and Santa Maria flight information regions (FIRs) over the North Atlantic, have the lowest traffic densities. Implementing navigational avoidance in such airspaces could lead to a higher success rate and overall mitigated warming due to lower congestion and higher warming impact. This could be part of a spatial targeting strategy, which aims to implement avoidance in the most warming locations whilst minimising disruption (Section 6 explores this in more detail).

Eastern European and oceanic airspaces have higher contrail warming at lower traffic levels

Traffic density (total flown distance distance per area as share of 2019 global mean) versus contrail warming per flown distance (share of 2019 global mean). Circle sizes proportional to contrail warming, darker colours stand for more contrail warming per flown distance.



Source: T&E (2025), based on Teoh et al. (2024) for the year 2019, re-run by ICL with an updated version of pycontrails (v0.54.8)

Figure 11: Graph from Transport and Environment's *How to Scale Up Contrail Avoidance in Europe* report [83] plotting contrail warming impact against traffic density across European and North Atlantic airspaces as percentages of the global average values in 2019.

However, around half of flights use the OTS, cruise in a narrow band of flight levels and fly in two unidirectional daily flows. Along the core tracks at peak times, local density is far higher than the FIR average suggests. This can be seen in Figure 12, visualising traffic over the Gander and Shanwick regions.



Figure 12: Visualisation of traffic in the North Atlantic produced by NATS

Meteorology adds a further constraint. Figure 13 shows that the contrail-forming share of North Atlantic airspace volume is typically 10-20 per cent in winter with peaks of 30 per cent, and is below 10 per cent in the summer. Airspace-based targeting that closes these regions could seriously complicate traffic management, especially if they happen to coincide with the OTS and cruise altitude bands. However, the strongest warming fraction of these contrails occupy roughly 5 per cent of airspace, and they are most prevalent in winter months when traffic is lowest. Hence, there is also the possibility of temporal targeting, where avoidance is implemented according to both daily and seasonal variations (see Section 6).

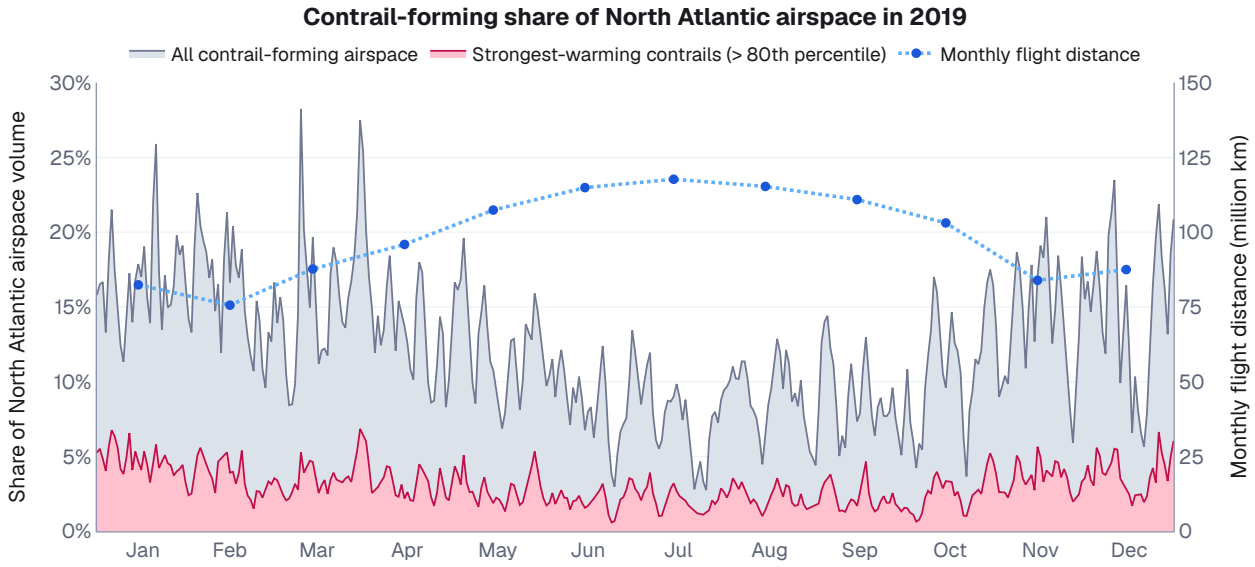


Figure 13: Percentage of airspace volume throughout 2019 that would form contrails of varying degrees of warming in the North Atlantic, superposed with monthly traffic volume in flight kilometers for the same year. Reproduced from Engerberg et al. (2025) [37] and Teoh et al. (2022b) [21].

The rerouting envelope is not fixed, and can be expanded with investment. Space-based ADS-B has lowered separation standards (longitudinal standards reduced from 40 nautical miles to 14 nautical miles; lateral standards reduced from 30 nautical miles to 19 nautical miles), allowed controllers to grant more level changes and started to reduce OTS reliance [32]. TBO concepts under development may further expand the envelope, but operational deployment is years away and will not eliminate the fundamental constraint of finite airspace capacity, which may be further challenged in the future by increased demand for commercial aviation.

5.5.5 Role in the dependency stack

→ **Depends on**: decision support systems (layer 3) for the avoidance options presented to planners and pilots.

→ **Informs**: regulation (layer 5), because what is operationally feasible defines what can reasonably be required.

Policy guardrail · Allow for variable operational flexibility

Because safety primacy and capacity constraints are permanent features of aviation, policy should accommodate variable operational envelopes rather than assume that every flight has the same flexibility. Performance-based targets that allow operators to optimise across their route networks are therefore likely to be more workable than flight-by-flight mandates. Avoidance may be infeasible for particular flights under conditions such as severe weather or high air-traffic density, and any performance framework should account explicitly for the CO_2 penalty of diversions. The relevant measure is net climate benefit across a portfolio, while ATM modernisation can expand the available operating envelope over time. The policy should nevertheless remain workable under current constraints.

5.6 Regulation and standards

This is the top of the stack: the mandates, incentives, international agreements and reporting and compliance mechanisms that depend on everything below.

5.6.1 Reporting, compliance and enforcement (MRV)

Investment-dependent | Structurally persistent

Any policy instrument generally requires a means of demonstrating compliance and assessing outcomes. MRV is a commonly used climate-policy framework. Its viability for contrail avoidance depends on the question it is designed to answer.

At any given moment, satellites detect approximately 46 per cent of the contrails present in the sky, across contrails of all ages; these contrails account for around 70 per cent of total contrail warming.

The remainder are too optically thin or too small to be detected, with the relevant observational constraints described in Section 5.2.2. An MRV system aiming to continuously track and verify every contrail produced by a flight would therefore be unable to account for roughly half of all contrails.

Designing a viable MRV system for contrails requires working within these structural limits. A more useful question is: if a contrail forms, how likely is it to be observed in a satellite image at some point during its lifetime? A contrail that is too optically thin or narrow to be seen in one image may become observable later. Accounting for this, the share of contrails observed at least once during their lifetime rises to 57 per cent at 2 km resolution and 64 per cent at 0.5 km resolution, and these contrails account for 97 per cent of total contrail warming [71]. These estimates use synthetic contrails generated by CoCiP against a clear sky, so they are upper limits.

This presents a more optimistic picture: if a contrail forms and contributes significantly to warming, current technology is likely to detect it at some point during its lifetime, although it may not be visible throughout. This could support a verification stage that assesses whether a contrail-mitigation strategy was successful without requiring continuous visibility. However, being observed is necessary but not sufficient for MRV: the contrail must also be attributed to the flight that produced it, and attribution depends on when the contrail is first seen.

A contrail can be confidently attributed to a specific flight only while it is young, when its geometry still traces the originating flight track and wind drift has not displaced it beyond the uncertainty region of advection models. Ground cameras observe contrails at or near the moment of formation, and the ground-camera literature reports attribution rates above 90 per cent [84]. However, ground cameras are concentrated over land and cannot provide the geographic coverage needed to scale up attribution, which points back to geostationary satellites. Satellites see contrails later, and attribution accuracy degrades rapidly with the delay between formation and first detection because of increasing uncertainty in wind speed over time [74]; matches beyond two hours are treated as erroneous by the same authors.

There are two metrics that describe attribution accuracy. Precision measures the share of contrail–flight matches that are correct, but it does not measure how many observed contrails are matched in total. An algorithm can achieve high precision by attributing only the clearest cases and leaving ambiguous contrails unmatched, so precision alone overstates how much of the contrail population attribution accounts for. For MRV purposes, recall (the fraction of observed contrails correctly attributed to their originating flight) matters as much as precision, because unattributed contrails are invisible to any compliance assessment and their omission may not be random across operators, routes or airspace density.

Figure 14 compares the precision and recall of contrail-to-flight attribution algorithms reported in several studies. Where a study reports these by contrail age, that variation is shown. Comprehensive data remains limited because this is an active field of research.

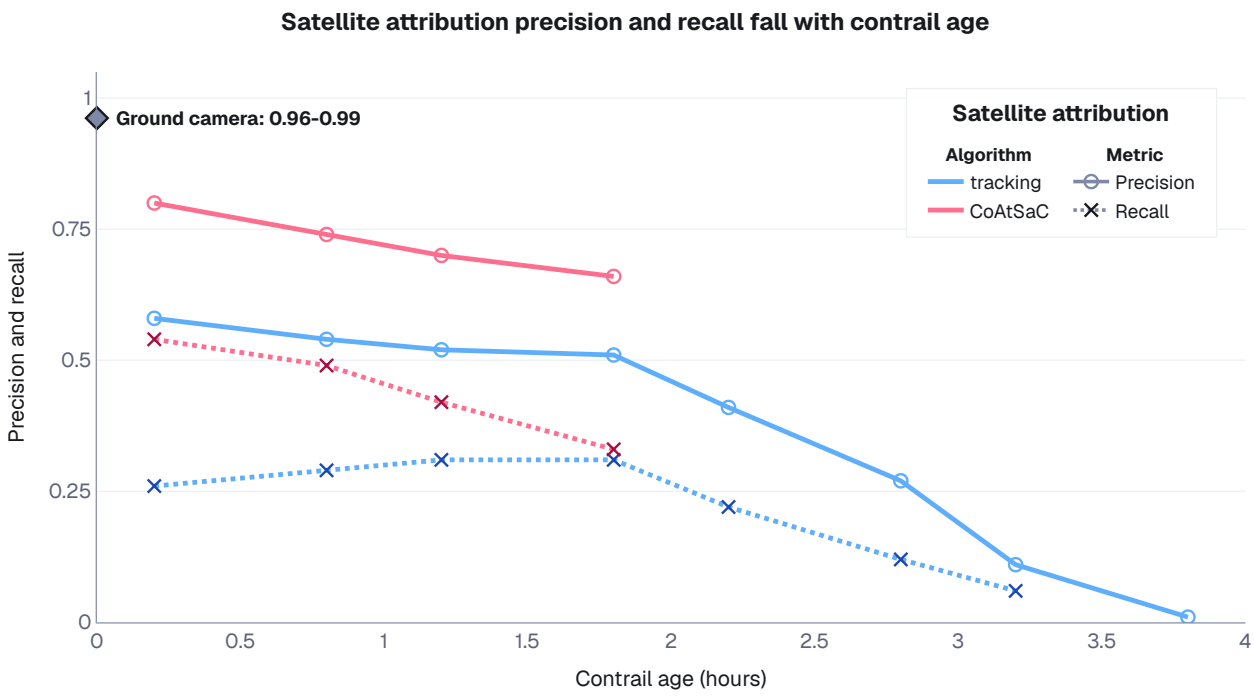


Figure 14: Accuracy of different attribution algorithms in recent studies. Satellite precision (solid) and recall (dotted) are per-contrail values for the tracking and CoAtSaC algorithms on the synthetic SynthOpenContrails benchmark, binned by contrail age and digitised from Sarna et al. (2025) [85]. The ground-camera point is a single aggregate attribution precision on new contrails, 0.962, from the Dalmau et al. (2026) [86]; the Nakoneczny et al. (2026) [87] reports a detection precision of 0.99. Metric definitions can differ between studies, so the ground-camera point gives context rather than a like-for-like comparison.

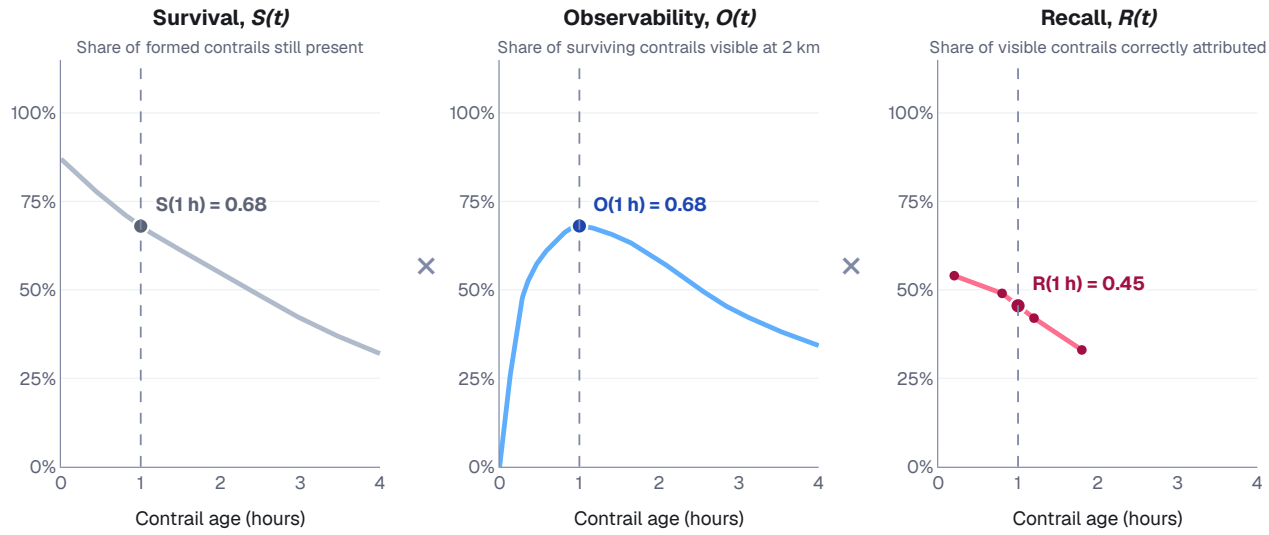
The attribution of an individual contrail depends on three variables, each a function of the contrail's age, t:

→ **Survival, $S(t)$:** the share of contrails formed that still exist at age t.

→ **Observability, $O(t)$:** the share of those still in existence that a satellite can detect, as shown in Figure 7.

→ **Recall, $R(t)$:** recall as defined above, which falls with contrail age as shown in Figure 14.

Each event is conditional on the one before, so their product, $S(t) \times O(t) \times R(t)$, gives the probability that a contrail is still present, visible and correctly attributed in a single satellite image taken at age t. Figure 15 shows this for an image taken one hour after formation with current 2 km satellites: 68 per cent of contrails are still present, 68 per cent of those are visible and 45 per cent of those are correctly attributed, so around 21 per cent of the contrails formed are correctly attributed in that image.



0.68 × 0.68 × 0.45 ≈ 0.21: about 21% of contrails formed are present, visible and correctly attributed in a single 2 km satellite image taken 1 hour after formation

Figure 15: Survival, observability at 2 km pixel resolution and CoAtSaC attribution recall by contrail age, each read off at one hour after formation (dashed line) and multiplied together. Observability is from Driver et al. (2025) [71] and recall from Sarna et al. (2025) [85]. Because the inputs come from separate studies, the product is illustrative.

This product is only a snapshot of one image at one age. A contrail needs to be attributed only once, and images taken at different ages largely capture the same contrails, so snapshots at different ages cannot be added together. As with observation above, the more useful question covers the whole lifetime: of the contrails observed at least once, how many are then correctly attributed? Attribution is normally attempted when a contrail is first detected, so the recall that matters is the recall at the age of first detection.

Figure 16 follows 100 contrails in this way under clear skies. With current 2 km satellite imagery, 57 are observed at least once in their lifetime, as described above. Their median first detection is 21 minutes after formation [71], when the best-performing satellite algorithm, CoAtSaC, correctly attributes around 53 per cent of contrails [85]. Around 30 of the 100 are therefore correctly attributed. At 0.5 km resolution, 64 are observed and the median first detection falls to 9 minutes, raising the number attributed to around 35. Higher resolution improves attribution in two ways: more contrails are observed at some point, and they are observed sooner, while recall is still high.

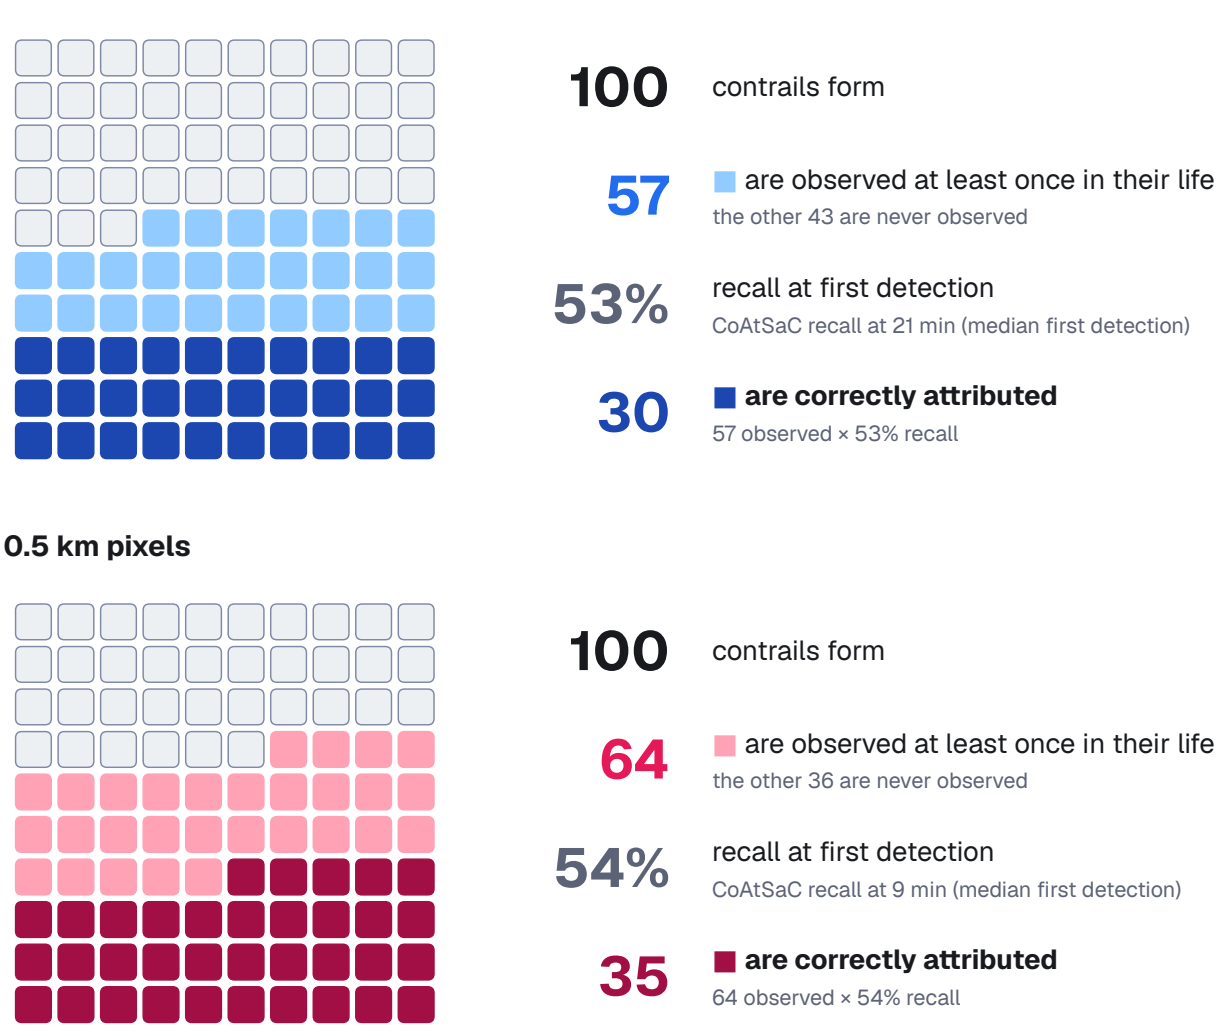


Figure 16: Estimated outcome for 100 contrails under clear skies at 2 km and 0.5 km satellite pixel resolution. The number observed at least once and the median time to first detection are from Driver et al. [71]; the number attributed applies CoAtSaC recall at that time [85]. Because the inputs come from separate studies, the estimates are illustrative.

Both estimates are idealised, as they assume clear skies and use attribution recall measured on a synthetic benchmark. As described in Section 5.2.2, clouds above a contrail can hide it from satellites entirely, and nearby contrails can obscure the linear shape that detection algorithms look for. In practice, these effects reduce the number of contrails observed and delay first detection. Under real-world conditions, Geraedts et al. found a mean first detection time of 41 minutes using 2 km Geostationary Operational Environmental Satellite 16 (GOES-16) imagery [74], by which time CoAtSaC recall has fallen to around 50 per cent. Figure 17 explores the same 100 contrails, with a delayed first detection time (based on the mean value reported in Geraedts et al.) and a reduced observability due to cloud occlusion. No published estimate yet shows how many contrails cloud hides over their whole lifetime, so we apply an illustrative 20 per cent penalty to observability. The number of contrails observed at least once drops to 46 out of 100, and the number correctly attributed drops from 30 to 23 (using a 2 km pixel resolution).

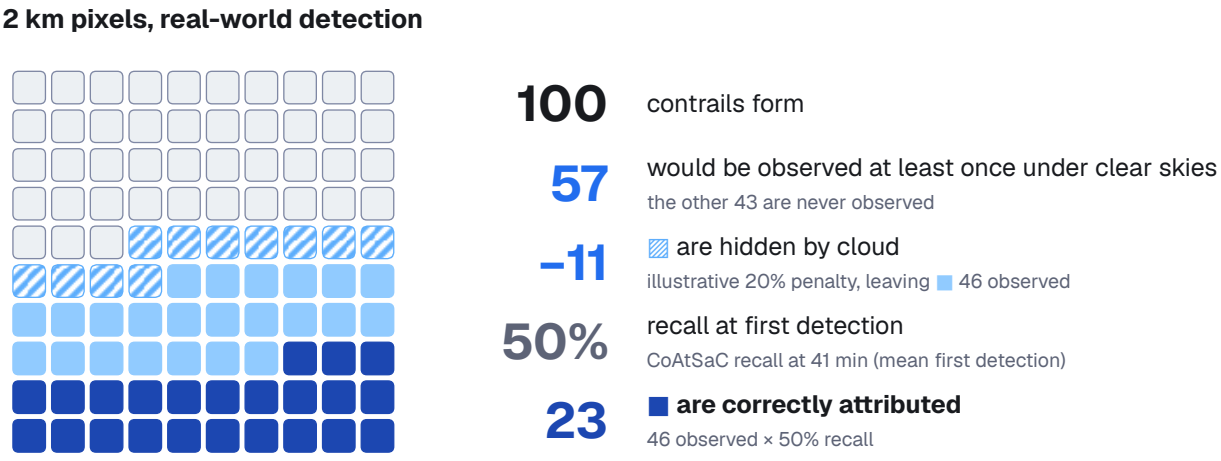


Figure 17: Estimated outcome for 100 contrails at 2 km satellite pixel resolution with illustrative cloud occlusion and with real-world detection timing. The real-world case applies CoAtSaC recall (Sarna et al., 2025) [85] at the 41-minute mean first detection time reported by Geraedts et al. [74], and an illustrative 20 per cent penalty on number of contrails observed once in their lifetime due to cloud cover.

This answers the more optimistic picture above. A contrail that contributes significantly to warming is likely to be seen at some point, but even under idealised conditions only around a third of contrails can be tied to the flight that produced it. Reliable individual-flight attribution at global scale would require detecting contrails within minutes of formation across all flight corridors. No current operational system achieves this, and even a fully realised low-Earth-orbit (LEO) constellation would not resolve the structural mismatch between optical detectability and attribution-compatible timing for roughly half of all contrails. Moreover, even after a contrail is detected, attribution can remain ambiguous in busy airspace, particularly because satellite images provide limited depth information.

MRV frameworks for contrail avoidance cannot be structured like carbon MRV, where emissions are directly measurable and attributable at source. The analysis above points towards model-based portfolio assessment as the primary viable architecture: operators report the contrail-related warming that qualifying models predict across their flights, and observational data validate model skill at the population level rather than verify individual flights. With appropriate model-selection and validation safeguards, assessing systematic patterns of avoidance behaviour can be more scalable and auditable than relying on flight-level observations of individual events, which may be disputed or selectively measured. Better satellites can improve model validation but do not change the fundamental architecture.

However, the credibility of a model-based approach depends on the models used and the rules governing their use. As noted in Section 5.4.1, two models assessing the same flight may produce materially different estimates. Without safeguards, model-based MRV could shift the risk of exploitation from selective measurement to selective modelling. The following safeguards are needed to address this:

Model selection. Allowing operators to report against any model would incentivise them to make a choice that credits them with the most avoided contrail warming at the lowest cost. Compliance claims should therefore be restricted to an approved range of models rather than left to operator choice. This does not have to be a single mandated model: a qualification standard, under which any model that meets a common benchmark against observations is eligible, preserves competition and room for methodological improvement while excluding models without demonstrated skill. Within a reporting period, the model version used should be fixed, so that claims cannot be re-run against a more favourable model after the fact. Where approved models disagree, claims could be reported against a common reference model, or with the spread across approved models stated. The qualifying standard should tighten as model skill improves, consistent with the phased approach set out in Section 5.4.

Validation against observations. Observational data (from satellites, ground cameras, in-flight humidity measurements and radiosondes) cannot verify individual flights at scale, but they can test whether a model's predictions hold up statistically across many flights. Standardised validation protocols must apply the same observational datasets, time periods and scoring methods to every candidate model in order to compare results between them. ContrailBench [80], described in Figure 9, is an early example of such a framework.

Forecast-skill thresholds. A minimum forecast-skill threshold sets the level below which compliance claims are not accepted. It must be carefully defined to prevent operators from meeting the threshold by optimising for a single score. As Figure 9 shows, a forecast's hit rate can be raised from 60 per cent to 93 per cent simply by expanding the forecast avoidance region, at the cost of flagging 24 per cent of flight distance for avoidance when observations indicate that contrails form over only 10 per cent of flight distance. A hit-rate threshold alone would therefore reward overprediction. Thresholds should instead assess hit rate jointly with the fraction of flight distance flagged, or compare the forecast rate with the observed rate of contrail formation.

Transparency of assumptions and inputs. Differences between models are often driven by assumptions about ice nucleation, contrail spreading, radiative transfer and climate metrics, as well as by input data such as weather forecasts, aircraft performance and fuel composition. These assumptions and inputs should be disclosed so that regulators and auditors can identify why models disagree. Common reporting formats would also support the interoperability described in Section 5.4.4. Input data matter as much as model physics: fuel-composition assumptions based on certificates rather than physical uplift are one known source of error (see Section 4.4.3).

Independent audit. Model qualification, validation results and operators' application of approved models to their flight data should be audited by a party independent of both model developers and operators. In a UK context, this role could sit with the regulator or a body it designates; Section 6.6 sets out one possible allocation.

Policy guardrail · Base MRV on model-based portfolio assessment using approved models

Individual-flight contrail attribution is structurally constrained: geostationary satellites often detect contrails after the attribution window has substantially narrowed, and no near-term observational system is expected to resolve this for the global fleet. MRV frameworks should therefore centre on model-based portfolio assessment, validated through statistical agreement with observations. Compliance claims should be restricted to an approved range of models that meet a common qualification standard, with model versions fixed for each reporting period. Supporting safeguards include standardised validation protocols, forecast-skill thresholds that cannot be met through over-prediction alone, transparency about model assumptions and input data and independent audit. The qualification standard should tighten as model skill improves.

5.6.2 International agreements and UK governance

Structurally persistent

The UK operates within the international aviation governance framework — the Chicago Convention, ICAO standards and recommended practices and bilateral air service agreements — but retains considerable domestic regulatory autonomy through the Civil Aviation Authority (CAA) and the Department for Transport (DfT). The international governance architecture, and the distribution of responsibility across actors, are structurally persistent. However, the governance constraint for UK contrail policy is not primarily about international consensus, but about designing domestic measures that are compatible with ICAO obligations and do not create undue competitive distortion for UK-based operators.

Several features of this landscape are worth noting:

Distributed responsibility. Even within UK airspace, contrail outcomes depend on decisions by airlines, NATS, the Met Office and the CAA. No single actor controls the full chain: an airline cannot reroute if NATS denies the altitude change, and NATS cannot mandate climate-optimal routing without airline cooperation on fuel costs. Individual actors will carry specific obligations, but someone must hold the whole chain together. An effective domestic policy should therefore name a single coordinating owner, such as the DfT or the CAA, to hold the delivery mandate. Section 6.6 sets out one possible allocation.

Compatibility with ICAO. The UK can act ahead of ICAO consensus (as it has on carbon pricing via the UK Emissions Trading Scheme (ETS)), but any contrail-specific measures would need to be framed consistently with ICAO's broader environmental work programme. This likely means performance-

based rather than prescriptive approaches, and voluntary or incentive-based mechanisms in the first instance, potentially transitioning to mandatory frameworks as international norms develop.

Regulatory design under uncertainty. Given the scientific and MRV constraints described above, prescriptive requirements risk becoming outdated quickly, while performance-based frameworks require MRV infrastructure that does not yet fully exist. This suggests an adaptive regulatory approach: setting frameworks that can tighten as capability matures, rather than locking in specific technical requirements at the outset.

Policy guardrail · Build an adaptive, ICAO-compatible framework

For the UK, the governance challenge is to design an adaptive domestic framework that remains compatible with ICAO while allowing the country to contribute to emerging international practice. Voluntary or incentive-based mechanisms may provide an appropriate starting point, with a possible transition towards mandatory obligations as MRV credibility and operational capability develop. This sequencing would allow experience to inform later requirements without precluding earlier action.

Policy guardrail · Name a single coordinating owner

No single actor controls the chain from forecast to flown trajectory: airlines, NATS, the Met Office and the CAA each control one step. Without a named owner, each actor can wait on the others. A single coordinating owner, such as the DfT or the CAA, should hold the mandate for delivery, setting escalation milestones and resolving cross-actor issues, while operational obligations remain with the actors best placed to carry them.

5.6.3 Role in the dependency stack

→ **Depends on**: operational procedures (layer 4) for what is feasible, and on MRV credibility (Section 5.6.1) for what can be enforced.

→ **Informs**: every layer below, through standards, funding and the demand signal that justifies investment.

Is cost a constraint?

Contrails.org estimates the global cost of contrail management at roughly $785 million per year, or about $21 per flight ($6.69–53.03) [88]. The fuel penalty from rerouting accounts for around 93 per cent of the estimated cost. Observation infrastructure (~$20 million per year), research and development (R&D; ~$18 million per year), additional staff (~$15 million per year) and forecasting (less than $1 million per year) make up the rest; MRV and regulatory costs are not included. The estimated costs of addressing investment-dependent constraints in layers 1–3 are comparatively modest, while the overall cost depends on how efficiently flights can be rerouted within the operational constraints of layer 4.

Even at the upper end of the range ($2 billion per year), navigational avoidance would cost at most 0.2 per cent of total airline expenditure ($900 billion per year [89]). It is one of the cheapest climate mitigation measures available in aviation, and cost should not be considered a barrier to policy development.

5.7 Summary

5.7.1 Constraints by layer and category

Layer	Investment-dependent	Structurally persistent
1. Observation	Humidity sensing; oceanic coverage; sensor data quality; satellite and machine-learning detection improvements at the margins	Optical-thickness limits on passive satellite detection
2. Meteorological data	Higher-resolution NWP; upper-tropospheric data assimilation; tailored forecast products and data access	Small-scale humidity variability below observational and model resolution
3. Decision support	Contrail-aware flight-planning tools; model validation; interoperability standards; support for smaller operators	Irreducible flight-level prediction uncertainty; choice of climate metric
4. Operational procedures	ATM modernisation, including Trajectory-Based Operations	Safety primacy; finite airspace capacity; tactical workload; variable feasibility and cost of avoidance
5. Regulation	MRV infrastructure; standards; incentive and policy design	Attribution versus detectability; distributed responsibility; ICAO governance

Table 5: Investment-dependent and structurally persistent constraints on navigational avoidance, by system layer.

5.7.2 What investment can fix

Meteorological observation gaps, infrastructure and capability immaturity and the limits of current technology are the primary addressable bottlenecks. These can be addressed through targeted funding, standards development, phased deployment programmes and continued technology improvement. Policy that fails to invest in these areas risks leaving contrail avoidance permanently imprecise.

5.7.3 What must be designed around

Fundamental atmospheric chaos at flight level, ATM safety primacy, the physics of contrail attribution and international governance architecture are persistent features of the problem. No feasible investment or technological breakthrough will eliminate them. Deferring action until ‘the science is ready’ or ‘governance catches up’ is therefore likely to delay action indefinitely.

6 Policy scope: targeted or universal?

Section 4 established what the different contrail mitigation mechanisms are, with pre-tactical navigational avoidance identified as the most promising near-term avenue. Section 5 explored the technical constraints on navigational avoidance and how they shape its practicability. The conclusion was that a viable policy must be portfolio-based, performance-based and phased. This section addresses the remaining questions: which flights should be covered, and which actors should carry which obligations? The analysis focuses on navigational avoidance; other mechanisms are noted only where the answer differs.

6.1 The efficiency argument

The argument for targeting rests on a fundamental feature of how contrails form. Contrail warming is not evenly distributed across the flight population. A small fraction of flights — those crossing ice-supersaturated regions at the right altitude, at night or in certain seasons and latitude bands — account for a disproportionate share of total contrail radiative forcing. Modelling studies consistently suggest that 2-5 per cent of flights are responsible for 80–90 per cent of contrail warming [19]. In contrast, around 30 per cent of flights form persistent contrails. A policy that imposes avoidance obligations uniformly across all flights therefore incurs a large share of the operational and economic cost of avoidance while capturing the climate benefit from only the targeted fraction. Targeting can, in principle, deliver most of the benefit at a fraction of the cost. However, it does not follow automatically that targeted policy is superior to universal policy. Targeting introduces its own costs: added policy complexity, operational feasibility challenges and inequities in burden distribution. Whether those costs are worth bearing depends substantially on which mechanism is deployed [90].

6.2 Taxonomy of targeting dimensions

Targeting can be applied across several distinct dimensions, each with different data requirements, policy instruments and operational implications, summarised in the table below.

Targeting dimension	What is targeted	Example policy instruments	Data requirement
Route-based	Specific origin-destination pairs with systematically high warming probability	Route-specific avoidance obligations	Historical contrail formation and warming estimates
Airspace-based	Defined geographic zones where persistent contrails are more likely	Zonal rules or conditional altitude obligations	Spatial contrail frequency maps
Temporal	Time periods with heightened warming	Time-of-day or seasonal avoidance obligations	Radiative forcing diurnal models (granular enough

Targeting dimension	What is targeted	Example policy instruments	Data requirement
	contribution (nighttime, winter seasons)		to observe daily fluctuations)
Forecast-driven	Individual flights predicted to encounter ice-supersaturated conditions	Dynamic, real-time obligations triggered by forecasts	Operational weather forecast skill

Table 6: Summary of the available targeting dimensions and corresponding data requirements and example policy instruments.

These dimensions are not mutually exclusive. The targeting problem is therefore not binary: it is a design choice about how many dimensions to target simultaneously and with what degree of precision.

The four dimensions also differ materially in their practical tractability. Route-based and airspace-based targeting share similar data requirements and face similar boundary-definition problems. Temporal targeting requires no real-time data and is stable year-to-year, which makes it the simplest to specify. Forecast-driven targeting is the most precise but requires operational meteorological infrastructure and introduces flight-by-flight variability in compliance obligations. These differences matter for policy design and are examined further below.

6.3 Costs of regulatory targeting

The efficiency gains from targeting are genuine, but they come with a set of costs that a universal approach avoids.

6.3.1 Policy complexity

A universal obligation is simple to specify: all flights, all the time. A targeted obligation requires precise definitions of what falls inside and outside the scope of the policy. Route-based targeting requires agreement on which routes qualify and borderline routes will create compliance disputes. Forecast-driven targeting requires specifying which model outputs trigger obligations, which models count and who bears liability when a forecast was wrong and avoidance was either attempted unnecessarily or omitted when it should have occurred. Model specification is in fact a requirement for all targeting approaches but for forecast-driven targeting the specification must be operationally live and carries direct flight-level liability, making the problem qualitatively more demanding.

The complexity compounds at the enforcement layer. Regulators need to verify that targeted operators complied and that non-targeted operators genuinely fell outside scope, not that they reclassified their operations to avoid falling within it. This creates a need for monitoring, reporting and verification infrastructure that is in practice as demanding as the infrastructure needed for universal policy, because the counterfactual (what the operator would have done under a universal obligation) still needs to be established.

6.3.2 Operational feasibility

Targeting introduces asymmetries between operators. Airlines operating predominantly high-impact routes bear a disproportionate compliance burden relative to competitors whose route networks fall largely outside the targeted set. This may reflect a rational climate-benefit allocation, but it creates competitive distortions that carriers on targeted routes will resist. If an obligation falls on flights crossing the North Atlantic at night, carriers whose schedules concentrate on those routes are systematically disadvantaged relative to carriers whose equivalent demand is met at different times or on different routes.

Forecast-driven targeting introduces additional operational asymmetry. Under a dynamic obligation, the avoidance requirement on any given flight is known only once a forecast is available, typically hours before departure. This compresses the operational window for planning and creates uncertainty for crews, dispatchers and slot planners. Airlines with limited meteorological access or short planning horizons may find dynamic targeting disproportionately burdensome compared to large carriers with dedicated flight-operations meteorology teams. Navigational avoidance is examined separately below because it is inherently conditional on atmospheric data.

6.3.3 Equity and legitimacy

Competitive asymmetry and procedural legitimacy are distinct concerns that targeted schemes must address separately.

On equity: Two airlines operating similar-impact flights should face similar obligations. Targeting based on route or airspace creates systematic winners and losers whose differential treatment reflects geography and schedule rather than any deliberate choice to cause greater harm. A carrier restructuring its transatlantic night programme because that route is high-impact is being treated differently from a carrier whose equivalent programme departs at a lower-impact time. That difference may not feel proportionate to those bearing the cost and it creates an incentive to compete on schedule design rather than on climate performance. This is among the strongest practical objections to route-based targeting and warrants explicit treatment in any framework that adopts it.

On legitimacy: Criteria for which routes or zones fall within scope must be technically defensible, transparent and subject to periodic review as the underlying science develops. Criteria that appear arbitrary, or that happen to disadvantage particular carriers without scientific justification, will be litigated. Frameworks should specify the data sources and models used to determine scope, publish them and establish a review cadence tied to improvements in contrail science.

6.4 Where targeting sits for navigational avoidance

Navigational avoidance has an implicit targeting logic built in. Deviations from the planned route occur when forecasts indicate an ice-supersaturated region ahead. An operator following a contrail-aware routing algorithm is already engaging in forecast-driven targeting: it avoids when the forecast says to avoid, and flies the standard route otherwise. Navigational avoidance is never truly universal as it is always conditional on atmospheric conditions.

The policy design question for navigational avoidance is therefore not whether to target at all, but at what level to impose the obligation. A universal obligation would require all operators to implement contrail-aware routing systems and follow them whenever the forecast indicates a high-probability warming region. A targeted obligation might restrict the requirement to specific route classes or airspace zones where contrail related warming is most severe. Given the operational and monitoring constraints discussed in Section 5, a universal but performance-based obligation (all operators, measured at portfolio level) may be more defensible than a targeted obligation that tries to define scope at flight or route level. The targeting is then delivered operationally (the algorithm targets the high-impact flights) rather than through the regulatory definition.

Other mechanisms

Scheduling avoidance: A night-time restriction is already a simple temporal targeting instrument that needs no forecasting. Adding route or airspace dimensions would make it harder to specify and enforce without a proportionate gain in precision.

Fuel composition: Unlike navigational avoidance, SAF has no credible role as a scalable contrail mitigation measure (Section 4.4). Thin supply would make targeting essential rather than an operational decision (Section 4.4.2). Most of the caveats regarding the complexity of targetting in this section will apply to SAF.

6.5 Is targeting worth it?

For navigational avoidance, yes, but through the avoidance algorithm rather than the regulatory definition. Operational targeting captures the efficiency gains of focusing on high-warming flights while avoiding the complexity, equity and legitimacy problems of scoping obligations by route or airspace.

The more important design consideration is that a narrow scope should not reduce policy ambition by default. A tightly scoped obligation may be administratively convenient, but it will deliver meaningful climate outcomes only if it includes an appropriate share of the high-impact, high-congestion routes where the problem is concentrated. Any targeted framework should therefore demonstrate, in quantitative terms, what share of total contrail radiative forcing its scope captures.

Policy guardrail · Apply to all operators equally, target by forecast

The efficiency case for targeting is genuine and should inform policy design, but the mechanism through which targeting is delivered matters as much as whether it is applied. For navigational avoidance, operational targeting embedded in the avoidance algorithm is preferable to regulatory targeting by route or airspace scope: the obligation should cover all operators, while the algorithm selects the high-impact flights. Where a policy framework defines scope or a phased rollout by route or airspace, it should set technically defensible criteria, provide for clear review and transition arrangements, use monitoring and verification requirements that distinguish genuine out-of-scope status from strategic reclassification and state what share of total contrail radiative forcing its scope captures, so that targeting is not used to reduce ambition.

6.6 Who carries which obligation

A universal obligation still has to be divided between the actors who control each step from forecast to flown trajectory (Section 5.6). The flight-planning phases in Section 4.6 give a natural way to do this, since each phase already has an established decision-maker. Table 7 sets out one possible allocation for navigational avoidance in the UK.

Phase	Actor	Possible obligation
Pre-tactical	Airlines and dispatchers	Plan with a qualified contrail-aware model; file plans that reflect it
Pre-tactical	Met Office	Provide ISSR forecast products to a defined standard
Network Manager review / tactical	NATS	Accommodate contrail-motivated requests where capacity allows; log refusals
Post-flight	Airlines	Report portfolio outcomes using existing debrief data
Oversight and coordination	CAA / DfT	Qualify models, audit outcomes and set escalation milestones; one body designated as the single coordinating owner

Table 7: Possible allocation of navigational avoidance obligations by flight-planning phase.

Most of these obligations extend an existing activity rather than create a new one: airlines already file software-optimised plans and already reconcile planned against flown routes in post-flight review (Section 8). Logging refusals matters for compliance: it separates flights where avoidance was infeasible because of ATM constraints from flights where it was never attempted, so operators are not penalised for capacity limits outside their control.

6.7 Summary

Targeted contrail avoidance is more efficient than universal avoidance, but efficiency is not the only consideration. Targeting by regulatory definition introduces policy complexity, competitive asymmetry and legitimacy questions that a universal obligation avoids, and a narrowly scoped obligation risks reducing ambition while preserving its appearance.

For navigational avoidance, these considerations point to a clear recommendation: universal coverage of all operators, performance-based obligations measured at portfolio level and targeting delivered operationally by contrail-aware routing that selects the high-impact flights, rather than by drawing regulatory boundaries around routes or airspace. Obligations within that framework fall to the actor who already controls each phase: airlines plan and report, the Met Office forecasts, NATS accommodates and logs, and the CAA and DfT qualify, audit and escalate.

7 Conclusion and recommendations

7.1 Conclusion

Persistent contrails are an important but currently under-managed part of aviation's climate impact. As research has developed, it has become increasingly clear that action must be taken despite the uncertainty naturally occurring in this area. The operational mechanisms required for contrail management seem straightforward in theory, but forecast error, airspace capacity, uncertain model outputs and uneven operational capability all limit the reduction that can be delivered in practice. The evidence does not support treating contrail management as a simple extension of existing carbon policy. Policy must be designed with these uncertainties in mind and evolve alongside global understanding and capability.

Taken together, the report's findings point towards a clear outcome. **Pre-tactical navigational avoidance is the strongest near-term policy pathway**, provided that it is implemented as a staged, performance-based system and remains subordinate to aviation safety and air traffic control. It can use existing aircraft, flight-planning processes and ATM infrastructure; it can target the small share of flights responsible for a disproportionate share of warming; and it can account for the fuel and operational cost of diversion within the same climate assessment.

Other mitigation mechanisms do not currently offer a better substitute. Scheduling avoidance may complement routing measures in specific circumstances, particularly where nighttime warming is material, but broad restrictions would cause greater disruption and could increase congestion and CO_2 emissions. SAF and other fuel interventions remain relevant to aviation decarbonisation, but their contrail effect is not sufficiently established or scalable to support near-term contrail crediting. Aircraft and engine modifications are valuable research directions, but certification and fleet turnover make them long-term options. The policy implication is therefore not to wait for a single technological solution, but to act on the most deployable mechanism while investing in the evidence and infrastructure needed to improve it.

7.2 Recommendations

1. **Safety comes first.** All contrail-avoidance actions must remain subordinate to ATC instructions and operational safety requirements, and safety-critical instructions always take precedence over climate-motivated routing (Section 2.2). Because safety and capacity limits vary from flight to flight, policy should allow for flights where avoidance is infeasible rather than penalise them (Section 5.5).
2. **Lead with pre-tactical navigational avoidance.** It is the only mechanism that combines near-term deployability, meaningful climate benefit and compatibility with existing ATM infrastructure (Section 4.6).
3. **Use portfolio-level monitoring and reporting.** Set fleet- or route-level performance targets, not per-flight targets. Individual-flight contrail prediction has an irreducible uncertainty floor (Section 5.4.1);

policy should not demand a level of attribution that the science cannot supply, regardless of investment.

4. **Avoid reliance on SAF or future aircraft technologies in the near term.** Exclude SAF and aircraft design modifications from near-term contrail crediting or compliance pathways. The science is unresolved, supply is inadequate and book-and-claim accounting is structurally incompatible with flight-specific contrail modelling (Section 4.4.4). Both remain worthwhile as parallel, long-term investments, not substitutes for operational action.
5. **Invest in infrastructure and weather data.** Treat meteorological and infrastructure investment as an explicit, funded workstream. Humidity sensing, higher-resolution NWP and interoperable flight-planning standards should be readiness milestones that gate any escalation from voluntary to mandatory obligations (Section 5.2 and Section 5.3).
6. **Do not conflate CO_2 and non-CO_2 effects.** The exchange rate between contrail warming and CO_2 depends on the metric chosen and remains contested, so contrail obligations should not be met with CO_2 reductions or allowances, nor the reverse (Section 5.4.2). Flight planning must still weigh the two, since avoidance burns extra fuel, but the choice of metric mainly affects marginal flights rather than the core case for avoidance.
7. **Name a single coordinating owner.** No single actor controls the chain from forecast to flown trajectory, so one body, such as the DfT or the CAA, should hold the mandate for delivery, setting escalation milestones and resolving cross-actor issues (Section 5.6). Operational obligations remain with the actors best placed to carry them (Section 6.6).
8. **Adopt adaptive, ICAO-compatible regulation.** The UK can move ahead of multilateral consensus, as it has with carbon pricing, provided obligations begin as voluntary or incentive-based measures and tighten as MRV credibility and operational capability develop (Section 5.6).
9. **Apply to all operators equally, target by forecast.** Structure navigational avoidance as a universal, performance-based obligation rather than a route- or airspace-scoped mandate, so that forecasts and routing software select the high-impact flights rather than a contestable regulatory boundary (Section 6.5). Where a scope or phased rollout is defined geographically, it should be technically justified, time-limited and state what share of total contrail warming it captures.

These recommendations are deliberately conditional. A policy should expand only when trials demonstrate that forecast skill, operational integration and measurement are good enough to support the next level of obligation. This creates a practical route from research and voluntary participation to enforceable performance requirements without treating uncertain model outputs as precise measurements.

7.3 Policy selection

The following checklist restates the three-part test set out in Section 2.2 against the architecture recommended above. It is a working tool for evaluating specific proposals based on the science and current operational experience of the sector. Any policy package should be assessed against these criteria when evaluating its effectiveness.

7.3.1 Safe

- ☐ Avoidance manoeuvres are explicitly subordinate to ATC/ATM clearance, and safety-critical instructions take precedence over any climate-motivated routing requirement.
- ☐ The proposal is viable under today's ATM infrastructure and does not treat future modernisation (e.g. TBO, free-route airspace) as a precondition for delivery.
- ☐ Override mechanisms are specified: avoidance obligations suspend immediately where safety demands it, and no operator, controller or crew is penalised for declining a manoeuvre on safety or capacity grounds.
- ☐ Avoidance manoeuvres comply with ICAO separation standards and applicable national airspace regulations, and do not introduce new hazards such as routing into known icing, turbulence or restricted airspace.
- ☐ Where the proposal alters approved flight plans or standard departure/arrival procedures, the required regulatory approval pathway is identified.

7.3.2 Effective

Net reduction in warming

- ☐ The proposal can demonstrate, through modelling or trial evidence, that its avoidance actions produce a net reduction in energy forcing rather than displacing warming to other regions, altitudes or times.
- ☐ The additional CO_2 from rerouting, delay or increased fuel burn is accounted for within the same climate metric as the avoided contrail forcing, so the policy cannot inadvertently increase total forcing.
- ☐ The climate benefit is assessed at whole-system level, including any second-order effects on congestion, daytime capacity or fleet deployment.

High coverage

- ☐ The proposal states what share of total contrail energy forcing its scope captures, in quantitative terms, and justifies that scope against the concentration of warming in roughly 2–5 per cent of flights.

- [ ] Where scope is defined by route, airspace, time or operator class, the proposal demonstrates that it does not systematically exclude the high-impact, high-congestion corridors where contrail forcing is concentrated.
- [ ] Where targeting is delivered operationally (by contrail models inside flight planning) rather than by regulatory definition, the proposal specifies which models qualify and what forecast skill threshold they must meet.

High compliance

- [ ] The compliance test is set at fleet, route or portfolio level and assessed through validated model outputs, not through verification of individual flights.
- [ ] Operators subject to the obligation have access to the real-time atmospheric guidance and decision-support tooling needed to comply, or a funded pathway to acquire it.
- [ ] The notice period for avoidance decisions is operationally realistic given flight planning and ATC coordination timescales, typically 12–24 hours ahead of departure for pre-tactical measures.
- [ ] The proposal specifies the data formats, climate metrics and reporting protocols to be used, or identifies which standards body is expected to produce them and by when.
- [ ] Model validation protocols, transparency requirements for model assumptions and inputs and a route to independent audit are defined.

Low disruption

- [ ] The proposal prefers pre-tactical routing adjustment over tactical rerouting, delay or cancellation wherever avoidance is technically feasible.
- [ ] It does not require avoidance on flights where the rerouting cost would outweigh the climate benefit.
- [ ] The distribution of operational cost across operators is assessed, including carriers whose networks concentrate on high-impact routes such as North Atlantic night services, and a mechanism for addressing disproportionate burden is identified.
- [ ] Smaller operators and less-resourced states are not excluded by disproportionate capability requirements; technical assistance or standardised tooling is available to them.

7.3.3 Feasible

Meteorological data

- [ ] The proposal does not assume forecast precision that current observation infrastructure cannot deliver.
- [ ] Where higher forecast accuracy is required, obligations are phased and gated against demonstrated improvements in forecast skill rather than fixed calendar dates.

- [ ] The proposal identifies a mechanism to fund or mandate the humidity sensing and model development investment that unlocks higher-accuracy avoidance, rather than assuming it will arrive on its own.

Infrastructure and capability

- [ ] Airlines, ANSPs and meteorological agencies subject to the obligation have, or have a funded pathway to acquire, the contrail-aware flight-planning systems, data feeds and procedures needed to comply.
- [ ] Readiness milestones are defined that gate escalation from voluntary or incentive-based measures to mandatory obligations.
- [ ] Interoperability standards are in place, or are being developed concurrently, so that different organisations assessing the same flight reach consistent conclusions.

Scientific and modelling uncertainty

- [ ] The proposal uses probabilistic metrics, such as expected energy forcing reduction across a portfolio, rather than deterministic per-flight measurement.
- [ ] Model outputs are treated as risk-weighted decision inputs rather than certified measurements, and the proposal states a minimum forecast skill threshold below which compliance claims are not accepted.
- [ ] The proposal distinguishes between uncertainty that model improvement will narrow and uncertainty that is structurally irreducible, and does not defer action pending the latter.

Operational and air traffic management

- [ ] The proposal accounts for the variable rerouting envelope available to different flights and does not assume all flights can achieve avoidance manoeuvres of the same magnitude.
- [ ] Avoidance is recognised as infeasible for some flights on some days, and operators are not penalised for genuine ATM constraints such as traffic density, airspace complexity or military restrictions.
- [ ] A single coordinating owner, such as the DfT or the CAA, is named, and coordination responsibilities across airlines, NATS, the Met Office and the CAA are specified, given that no single actor controls the full chain from forecast to executed trajectory (Section 6.6).

7.4 Closing statement

The core requirement for contrail policy is not whether it can predict every contrail perfectly. No current system can. The requirement is whether the policy can produce a reliable net reduction in warming while remaining safe, operationally workable and accountable for uncertainty. A phased navigational-avoidance framework meets that test more closely than the alternatives available today.

The next step is therefore to build the capability to act: run representative flight trials, improve atmospheric observations and forecasts, establish common data and reporting standards and measure outcomes at portfolio level. Those steps should be linked to explicit readiness thresholds and reviewed against realised climate benefit, fuel penalty, disruption and equity across operators. This approach preserves the value of early action without locking policy into assumptions that later evidence may overturn.

Contrail management should become part of aviation climate policy now, but it should do so as an adaptive operational programme rather than a one-off mandate. Acting on that basis allows the UK to capture near-term climate benefits, learn from implementation and contribute a credible model for wider international adoption.

8 Appendix

8.1 The flight-planning process

Effective contrail management requires interventions that fit within the existing structures of aviation operations. This section describes the four principal phases of the commercial flight-planning process, from strategic network scheduling through to post-flight review, to establish the operational context within which any contrail avoidance policy would need to function.

8.1.1 Phase 1: strategic planning

The flight-planning process can begin weeks or even months before an aircraft departs. Airlines plan their route networks based on a combination of factors, including passenger demand, available fleet capacity, commercial strategy and the availability of airport slots. Airport slot coordinators play a key role at this stage, allocating the take-off and landing windows that airlines are permitted to use.

By the end of the strategic planning phase, an airline will have established which flights it intends to operate, their origin and destination airports and their scheduled departure and arrival times. However, the specific route each flight will take through the airspace has not yet been determined.

8.1.2 Phase 2: pre-tactical planning

In the pre-tactical planning phase, airline dispatchers use specialist software to generate a digital flight plan in a standardised industry format. This plan is typically optimised across several parameters, including fuel burn, flight duration and overall operating cost.

Dispatchers must also consult a range of published sources when preparing the flight plan, including the Air Traffic Flow and Capacity Management (ATFCM) plan and aviation weather forecasts. These inform decisions about routing and altitude to ensure the flight can be conducted safely and efficiently.

The completed flight plan must be submitted for approval at least one hour before the scheduled departure time, though it is common practice to file flight plans a full day in advance. Approval is granted through EUROCONTROL's IFPS. Once accepted, the flight plan is automatically distributed to all relevant ATC units along the planned route [91].

Immediately prior to departure, the flight crew contacts the ATC tower at their departure airport to request take-off clearance. If approved, the aircraft is assigned a Calculated Take-Off Time (CTOT), which defines a tolerance window of minus five to plus ten minutes within which the aircraft must depart.

8.1.3 Phase 3: tactical (in-flight) management

Once airborne, both the flight crew and ATC continuously monitor the flight's progress and any changes in conditions. As the aircraft moves through different regions of airspace, the crew communicates with a succession of air traffic controllers, each responsible for a distinct sector. Hand-offs between controllers are coordinated to ensure continuity of service.

It is common for flights to deviate from their filed flight plan during this phase. Reasons for deviation may include adverse weather conditions, changes in fuel consumption, airspace restrictions or the need to maintain safe separation from other aircraft. In such cases, revised clearances are issued by ATC and the crew adjusts the flight path accordingly.

8.1.4 Phase 4: post-flight review

Following the completion of each flight, airlines carry out a standard review process. This typically involves reconciling the filed flight plan against the actual route flown, reviewing fuel burn data and documenting any delays or deviations. The outcomes of this analysis feed into operational improvements and future planning cycles.

References

[1] Aviation Impact Accelerator, “Searching for blue skies: Policy roadmap for navigational contrail avoidance,” Cambridge, UK, White Paper, Nov. 2026. [Online]. Available: https://aiazero.org/docs/policy-roadmap-for-navigational-contrail-avoidance.pdf

[2] U. Schumann, “Influence of propulsion efficiency on contrail formation,” *Aerospace Science and Technology*, vol. 4, no. 6, pp. 391–401, Sep. 2000, doi: 10.1016/S1270-9638(00)01062-2.

[3] R. Paoli, J. Hélie, and T. Poinsot, “Contrail formation in aircraft wakes,” *Journal of Fluid Mechanics*, vol. 502, pp. 361–373, Mar. 2004, doi: 10.1017/S0022112003007808.

[4] D. K. Singh, S. Sanyal, and D. J. Wuebbles, “Understanding the role of contrails and contrail cirrus in climate change: a global perspective,” *Atmospheric Chemistry and Physics*, vol. 24, no. 16, pp. 9219–9262, Aug. 2024, doi: 10.5194/acp-24-9219-2024.

[5] B. Kärcher, “Formation and radiative forcing of contrail cirrus,” *Nature Communications*, vol. 9, no. 1, p. 1824, May 2018, doi: 10.1038/s41467-018-04068-0.

[6] D. S. Lee *et al.*, “The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018,” *Atmospheric Environment*, vol. 244, p. 117834, Jan. 2021, doi: 10.1016/j.atmosenv.2020.117834.

[7] M. Klöwer, M. R. Allen, D. S. Lee, S. R. Proud, L. Gallagher, and A. Skowron, “Quantifying aviation's contribution to global warming,” *Environmental Research Letters*, vol. 16, no. 10, Oct. 2021, doi: 10.1088/1748-9326/ac286e.

[8] J. R. Smith *et al.*, “The climate opportunities and risks of contrail avoidance,” *Nature Communications*, vol. 17, no. 1, p. 2092, Jan. 2026, doi: 10.1038/s41467-026-68784-8.

[9] J. A. Knox, “Possible Mechanisms of Clear-Air Turbulence in Strongly Anticyclonic Flows,” *Monthly Weather Review*, vol. 125, no. 6, pp. 1251–1259, Jun. 1997, doi: 10.1175/1520-0493(1997)125<1251:PMOCAT>2.0.CO;2.

[10] “World Area Forecast Centre (WAFC).” Accessed: Jun. 16, 2026. [Online]. Available: https://www.metoffice.gov.uk/services/transport/aviation/regulated/international-aviation/wafc

[11] National Air Traffic Services, “Aeronautical Information Regulation and Control.” Accessed: Jul. 09, 2026. [Online]. Available: https://www.aurora.nats.co.uk/htmlAIP/Publications/2026-07-09-AIRAC/html/index-en-GB.html

[12] ICAO, “Annex 3 — Meteorological Service for International Air Navigation - 20th Ed..” Accessed: Jun. 16, 2026. [Online]. Available: https://elibrary.icao.int/reader/264207/&returnUrl%3DaHR0cHM6Ly9lbGlicmFyeS5pY2FvLmludC9wcm9kdWN0LzI2NDIwNw%3D%3D?productType=ebook

[13] P. Spichtinger, K. Gierens, and W. Read, “The global distribution of ice-supersaturated regions as seen by the Microwave Limb Sounder,” *Quarterly Journal of the Royal Meteorological Society*, vol. 129, no. 595, pp. 3391–3410, 2003, doi: 10.1256/qj.02.141.

[14] S. Marquart, M. Ponater, F. Mager, and R. Sausen, “Future Development of Contrail Cover, Optical Depth, and Radiative Forcing: Impacts of Increasing Air Traffic and Climate Change,” *Journal of Climate*, vol. 16, no. 17, pp. 2890–2904, Sep. 2003, doi: 10.1175/1520-0442(2003)016<2890:FDOCCO>2.0.CO;2.

[15] R. Dischl, S. Kaufmann, and C. Voigt, “Regional and Seasonal Dependence of the Potential Contrail Cover and the Potential Contrail Cirrus Cover over Europe,” *Aerospace*, vol. 9, no. 9, p. 485, Sep. 2022, doi: 10.3390/aerospace9090485.

[16] E. J. Jensen, A. S. Ackerman, D. E. Stevens, O. B. Toon, and P. Minnis, “Spreading and growth of contrails in a sheared environment,” *Journal of Geophysical Research: Atmospheres*, vol. 103, no. D24, pp. 31557–31567, 1998, doi: 10.1029/98JD02594.

[17] B. G. Liepert, “The physical concept of climate forcing,” *WIREs Climate Change*, vol. 1, no. 6, pp. 786–802, 2010, doi: 10.1002/wcc.75.

[18] R. Meerkötter, U. Schumann, D. R. Doelling, P. Minnis, T. Nakajima, and Y. Tsushima, “Radiative forcing by contrails,” *Annales Geophysicae*, vol. 17, no. 8, pp. 1080–1094, Aug. 1999, doi: 10.1007/s00585-999-1080-7.

[19] R. Teoh *et al.*, "Global aviation contrail climate effects from 2019 to 2021," *Atmospheric Chemistry and Physics*, vol. 24, no. 10, pp. 6071–6093, May 2024, doi: 10.5194/acp-24-6071-2024.

[20] M. Diao *et al.*, "Distributions of ice supersaturation and ice crystals from airborne observations in relation to upper tropospheric dynamical boundaries," *Journal of Geophysical Research: Atmospheres*, vol. 120, no. 10, pp. 5101–5121, 2015, doi: 10.1002/2015JD023139.

[21] R. Teoh *et al.*, "Aviation contrail climate effects in the North Atlantic from 2016 to 2021," *Atmospheric Chemistry and Physics*, vol. 22, no. 16, pp. 10919–10935, Aug. 2022, doi: 10.5194/acp-22-10919-2022.

[22] U. Schumann, "A contrail cirrus prediction model," *Geoscientific Model Development*, vol. 5, no. 3, pp. 543–580, May 2012, doi: 10.5194/gmd-5-543-2012.

[23] U. Schumann and A. Seifert, "On the Weather Impact of Contrails: New Insights from Coupled ICON–CoCiP Simulations," *Atmospheric Chemistry and Physics*, vol. 25, no. 24, pp. 18571–18597, Dec. 2025, doi: 10.5194/acp-25-18571-2025.

[24] V. R. Meijer, S. D. Eastham, I. A. Waitz, and S. R. H. Barrett, "Contrail altitude estimation using GOES-16 ABI data and deep learning," *Atmospheric Measurement Techniques*, vol. 17, no. 20, pp. 6145–6162, Oct. 2024, doi: 10.5194/amt-17-6145-2024.

[25] L. Wilhelm, K. Gierens, and S. Rohs, "Meteorological Conditions That Promote Persistent Contrails," *Applied Sciences*, vol. 12, no. 9, p. 4450, Jan. 2022, doi: 10.3390/app12094450.

[26] R. Ren, J. Zhong, D. Guo, R. Wang, and C. Hurter, "Review of ADS-B Data Usage with the focus on Data Cleaning," *Journal of Open Aviation Science*, vol. 4, no. 2, Apr. 2026, doi: 10.59490/joas.2026.8467.

[27] K. M. Gierens, L. Lim, and K. Eleftheratos, "A review of various strategies for contrail avoidance," *Open Atmospheric Science Journal*, vol. 2, pp. 1–7, Jan. 2008, Accessed: Jun. 22, 2026. [Online]. Available: http://www.bentham.org/open/toascj/openaccess2.htm

[28] A. M. Frias, M. L. Shapiro, Z. Engberg, R. Zopp, M. Soler, and M. E. J. Stettler, "Feasibility of contrail avoidance in a commercial flight planning system: an operational analysis," *Environmental Research: Infrastructure and Sustainability*, vol. 4, no. 1, Mar. 2024, doi: 10.1088/2634-4505/ad310c.

[29] R. S. Märkl *et al.*, "Powering aircraft with 100% sustainable aviation fuel reduces ice crystals in contrails," *Atmospheric Chemistry and Physics*, vol. 24, no. 6, pp. 3813–3837, Mar. 2024, doi: 10.5194/acp-24-3813-2024.

[30] T. R. Dean *et al.*, "Impact of forecast stability on navigational contrail avoidance," *Environmental Research: Infrastructure and Sustainability*, vol. 5, no. 4, p. 45008, Nov. 2025, doi: 10.1088/2634-4505/ae1da5.

[31] R. Sausen *et al.*, "Can we successfully avoid persistent contrails by small altitude adjustments of flights in the real world?," *Meteorologische Zeitschrift*, vol. 33, no. 1, pp. 83–98, Jun. 2024, doi: 10.1127/metz/2023/1157.

[32] J. Molloy *et al.*, "Design Principles for a Contrail-Minimizing Trial in the North Atlantic," *Aerospace*, vol. 9, no. 7, p. 375, Jul. 2022, doi: 10.3390/aerospace9070375.

[33] S. M. Hofer and K. M. Gierens, "Synoptic and microphysical lifetime constraints for contrails," *Atmospheric Chemistry and Physics*, vol. 25, no. 16, pp. 9235–9247, Aug. 2025, doi: 10.5194/acp-25-9235-2025.

[34] A. Sonabend-W *et al.*, "Feasibility test of per-flight contrail avoidance in commercial aviation," *Communications Engineering*, vol. 3, no. 1, Dec. 2024, doi: 10.1038/s44172-024-00329-7.

[35] J. Rosenow and M. Luo, "A Contrail Life Cycle Model with Interaction of Overlapping Contrails," *Aerospace*, vol. 13, no. 2, p. 164, Feb. 2026, doi: 10.3390/aerospace13020164.

[36] H. Heggarty, "Operation Blue Skies Takes Off: Landmark Oceanic Airspace Trial to Test AI-Driven Contrail Avoidance." Accessed: Sep. 18, 2026. [Online]. Available: https://www.nats.aero/news/operation-blue-skies-takes-off-landmark-oceanic-airspace-trial-to-test-ai-driven-contrail-avoidance/

[37] Z. Engberg, R. Teoh, T. Abbott, T. Dean, M. E. J. Stettler, and M. L. Shapiro, "Forecasting contrail climate forcing for flight planning and air traffic management applications: the CocipGrid model in pycontrails 0.51.0," *Geoscientific Model Development*, vol. 18, no. 2, pp. 253–286, Jan. 2025, doi: 10.5194/gmd-18-253-2025.

[38] T. Sankar *et al.*, "Efficacy of Scalable Airline-led Contrail Avoidance." Accessed: Sep. 18, 2026. [Online]. Available: http://arxiv.org/abs/2603.06909

[39] S. Andrews *et al.*, “Understanding Contrail Management: Opportunities, Challenges, and Insights.” Accessed: Jun. 22, 2026. [Online]. Available: https://rmi.org/insight/understanding-contrail-management-opportunities-challenges-and-insights?submitted=1#thank-you

[40] E. B. Peterson, K. Neels, N. Barczi, and T. Graham, “The Economic Cost of Airline Flight Delay,” *Journal of Transport Economics and Policy*, vol. 47, no. 1, pp. 107–121, 2013, Accessed: Jun. 22, 2026. [Online]. Available: https://www.jstor.org/stable/24396355

[41] N. Stuber, P. Forster, G. Rädel, and K. Shine, “The importance of the diurnal and annual cycle of air traffic for contrail radiative forcing,” *Nature*, vol. 441, no. 7095, pp. 864–867, Jun. 2006, doi: 10.1038/nature04877.

[42] Department for Transort, “Night flight restrictions: Heathrow, Gatwick and Stansted airports from October 2025.” Accessed: Jul. 10, 2026. [Online]. Available: https://www.gov.uk/government/consultations/night-flight-restrictions-heathrow-gatwick-and-stansted-airports-from-october-2025/night-flight-restrictions-heathrow-gatwick-and-stansted-airports-from-october-2025

[43] N. Stuber and P. Forster, “The impact of diurnal variations of air traffic on contrail radiative forcing,” *Atmospheric Chemistry and Physics*, vol. 7, no. 12, pp. 3153–3162, Jun. 2007, doi: 10.5194/acp-7-3153-2007.

[44] J. Rosenow, H. Fricke, and L. Sherry, “Time of the day-dependent impact of Contrail Avoidance Strategies on Airline Delay Costs,” in *Fifteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2023)*, Savannah, GA, USA, Jun. 2023.

[45] K. Kölker, Z. Zengerling, M. Kühlen, K. Lütjens, and F. Linke, “Assessing the impact of contrail avoidance through rescheduling on airline network flows: A case study of North Atlantic flights,” *Transportation Research Part A: Policy and Practice*, vol. 187, p. 104155, Sep. 2024, doi: 10.1016/j.tra.2024.104155.

[46] R. Teoh *et al.*, “Targeted Use of Sustainable Aviation Fuel to Maximize Climate Benefits,” *Environmental Science & Technology*, vol. 56, no. 23, pp. 17246–17255, Dec. 2022, doi: 10.1021/acs.est.2c05781.

[47] C. Voigt *et al.*, “Cleaner burning aviation fuels can reduce contrail cloudiness,” *Communications Earth & Environment*, vol. 2, no. 1, p. 114, Jun. 2021, doi: 10.1038/s43247-021-00174-y.

[48] R. S. Märkl *et al.*, “Powering aircraft with 100 % sustainable aviation fuel reduces ice crystals in contrails,” *Atmospheric Chemistry and Physics*, vol. 24, no. 6, pp. 3813–3837, Mar. 2024, doi: 10.5194/acp-24-3813-2024.

[49] F. Yu, B. Kärcher, and B. E. Anderson, “Revisiting Contrail Ice Formation: Impact of Primary Soot Particle Sizes and Contribution of Volatile Particles,” *Environmental Science & Technology*, vol. 58, no. 40, pp. 17650–17660, Oct. 2024, doi: 10.1021/acs.est.4c04340.

[50] C. Voigt *et al.*, “Substantial aircraft contrail formation at low soot emission levels,” *Nature*, vol. 652, no. 8108, pp. 112–118, Apr. 2026, doi: 10.1038/s41586-026-10286-0.

[51] “Sustainable Aviation Fuel (SAF) Mandate.” Accessed: Sep. 10, 2026. [Online]. Available: https://www.gov.uk/government/collections/sustainable-aviation-fuel-saf-mandate

[52] E. Woeldgen, R. Teoh, M. E. J. Stettler, and R. Malina, “Sustainable Aviation Fuel Deployment Strategies in Europe: Supply Chain Implications and Climate Benefits,” *Environmental Science & Technology*, vol. 59, no. 25, pp. 12447–12457, Jul. 2025, doi: 10.1021/acs.est.5c02364.

[53] G. Quante, S. Voß, N. Bullerdiek, C. Voigt, and M. Kaltschmitt, “Hydroprocessing of fossil fuel-based aviation kerosene – Technology options and climate impact mitigation potentials,” *Atmospheric Environment: X*, vol. 22, p. 100259, Apr. 2024, doi: 10.1016/j.aeaoa.2024.100259.

[54] Engine and Aircraft Certification Working Group, “Engine and Aircraft Certification Working Group (EACWG) report,” technical report, Jun. 2017. Accessed: Jul. 01, 2026. [Online]. Available: https://www.easa.europa.eu/en/document-library/general-publications/engine-and-aircraft-certification-working-group-eacwg-report

[55] “An overview of EASA’s work on aircraft certification.” Accessed: Jul. 01, 2026. [Online]. Available: https://www.easa.europa.eu/en/light/topics/idea-sky-how-easa-ensures-aircraft-safe-fly

[56] T. Arblaster, N. Thonemann, and B. Steubing, “Air traffic growth jeopardises European aviation’s climate mitigation efforts despite the substantial potential of hydrogen,” *Communications Earth & Environment*, vol. 6, no. 1, p. 976, Nov. 2025, doi: 10.1038/s43247-025-02935-5.

[57] K. Hunt, “Beyond carbon dioxide: Aviation needs a multi-pronged strategy to address contrails and reduce climate impacts.” Accessed: Jun. 29, 2026. [Online]. Available: https://www.catf.us/2025/03/beyond-carbon-dioxide-aviation-needs-multi-pronged-strategy-address-contrails-reduce-climate-impacts/

[58] J. Callard, C. Klumpp, D. Weintraub, S. Henninger, and P. Jeschke, “Reducing contrail formation by engine design and operating point adaptation of aircraft engines,” *CEAS Aeronautical Journal*, vol. 17, no. 2, pp. 913–924, Apr. 2026, doi: 10.1007/s13272-025-00887-2.

[59] X. Gao, A. Isoldi, D. Nalianda, and T. Nikolaidis, “Water extraction in aero gas turbines for contrail mitigation,” *The Aeronautical Journal*, vol. 128, no. 1325, pp. 1344–1361, Jul. 2024, doi: 10.1017/aer.2024.22.

[60] F. Noppel, R. Singh, and M. Taylor, “Novel engine concept to suppress contrail and cirrus cloud formation,” in *TAC (Thermal Methods Group)*, Jul. 2006.

[61] R. Pouzolz, O. Schmitz, and H. Klingels, “Evaluation of the Climate Impact Reduction Potential of the Water-Enhanced Turbofan (WET) Concept,” *Aerospace*, vol. 8, no. 3, p. 59, Mar. 2021, doi: 10.3390/aerospace8030059.

[62] “WET engine: MTU Aero Engines develops drive of the future with wet combustion.” Accessed: Jun. 29, 2026. [Online]. Available: https://www.mtu.de/newsroom/press/latest-press-releases/press-release-detail/wet-engine-mtu-aero-engines-develops-drive-of-the-future-with-wet-combustion/

[63] U. Ebner, “MTU verabschiedet sich von Wasserdampfeinspritzung.” Accessed: Jun. 29, 2026. [Online]. Available: https://www.flugrevue.de/flugzeugbau/wet-engine-mtu-verabschiedet-sich-von-wasserdampfeinspritzung/

[64] F. Haglind, “Potential of lowering the contrail formation of aircraft exhausts by engine re-design,” *Aerospace Science and Technology*, vol. 12, no. 6, pp. 490–497, Sep. 2008, doi: 10.1016/j.ast.2007.12.001.

[65] U. Burkhardt, L. Bock, and A. Bier, “Mitigating the contrail cirrus climate impact by reducing aircraft soot number emissions,” *npj Climate and Atmospheric Science*, vol. 1, no. 1, p. 37, Oct. 2018, doi: 10.1038/s41612-018-0046-4.

[66] M. A. H. Khan, J. Brierley, K. N. Tait, S. Bullock, D. E. Shallcross, and M. H. Lowenberg, “The Emissions of Water Vapour and NOx from Modelled Hydrogen-Fuelled Aircraft and the Impact of NOx Reduction on Climate Compared with Kerosene-Fuelled Aircraft,” *Atmosphere*, vol. 13, no. 10, p. 1660, Oct. 2022, doi: 10.3390/atmos13101660.

[67] A. Bier, S. Unterstrasser, J. Zink, D. Hillenbrand, T. Jurkat-Witschas, and A. Lottermoser, “Contrail formation on ambient aerosol particles for aircraft with hydrogen combustion: a box model trajectory study,” *Atmospheric Chemistry and Physics*, vol. 24, no. 4, pp. 2319–2344, Feb. 2024, doi: 10.5194/acp-24-2319-2024.

[68] A. Batra, R. Raute, and R. Camilleri, “On the Range Equation for Hybrid-Electric Aircraft,” *Aerospace*, vol. 10, no. 8, p. 687, Aug. 2023, doi: 10.3390/aerospace10080687.

[69] A. Petzold *et al.*, “Global-scale atmosphere monitoring by in-service aircraft – current achievements and future prospects of the European Research Infrastructure IAGOS,” *Tellus B: Chemical and Physical Meteorology*, vol. 67, no. 1, p. 28452, Dec. 2015, doi: 10.3402/tellusb.v67.28452.

[70] D. Atlas, Z. Wang, and D. P. Duda, “Contrails to Cirrus—Morphology, Microphysics, and Radiative Properties,” *Journal of Applied Meteorology and Climatology*, vol. 45, no. 1, pp. 5–19, Jan. 2006, doi: 10.1175/JAM2325.1.

[71] O. G. A. Driver, M. E. J. Stettler, and E. Gryspeerdt, “Factors limiting contrail detection in satellite imagery,” *Atmospheric Measurement Techniques*, vol. 18, no. 5, pp. 1115–1134, Mar. 2025, doi: 10.5194/amt-18-1115-2025.

[72] R. Chevallier, M. Shapiro, Z. Engberg, M. Soler, and D. Delahaye, “Linear Contrails Detection, Tracking and Matching with Aircraft Using Geostationary Satellite and Air Traffic Data,” *Aerospace*, vol. 10, no. 7, p. 578, Jul. 2023, doi: 10.3390/aerospace10070578.

[73] E. Gryspeerdt *et al.*, “Operational differences lead to longer lifetimes of satellite detectable contrails from more fuel efficient aircraft,” *Environmental Research Letters*, vol. 19, no. 8, p. 84059, Aug. 2024, doi: 10.1088/1748-9326/ad5b78.

[74] S. Geraedts *et al.*, “A scalable system to measure contrail formation on a per-flight basis,” *Environmental Research Communications*, vol. 6, no. 1, p. 15008, Jan. 2024, doi: 10.1088/2515-7620/ad11ab.

[75] K. Gierens and P. Spichtinger, “On the size distribution of ice-supersaturated regions in the upper troposphere and lowermost stratosphere,” *Annales Geophysicae*, vol. 18, no. 4, pp. 499–504, Apr. 2000, doi: 10.1007/s00585-000-0499-7.

[76] “Operational configurations of the Integrated Forecasting System (IFS) - Forecast User - ECMWF Confluence Wiki.” Accessed: Jun. 17, 2026. [Online]. Available: https://confluence.ecmwf.int/pages/viewpage.action?pageId=324860211

[77] P. Reutter, P. Neis, S. Rohs, and B. Sauvage, “Ice supersaturated regions: properties and validation of ERA-Interim reanalysis with IAGOS in situ water vapour measurements,” *Atmospheric Chemistry and Physics*, vol. 20, no. 2, pp. 787–804, Jan. 2020, doi: 10.5194/acp-20-787-2020.

[78] P. Hodgson, “Contrail Regions 101 (Part I).” Accessed: Sep. 09, 2026. [Online]. Available: https://notebook.contrails.org/contrail-regions-101-part-i/

[79] J. C. Platt *et al.*, “The effect of uncertainty in humidity and model parameters on the prediction of contrail energy forcing,” *Environmental Research Communications*, Aug. 2024, doi: 10.1088/2515-7620/ad6ee5.

[80] P. Hodgson and T. Abbott, “ContrailBench v1.” Accessed: Sep. 11, 2026. [Online]. Available: https://bench.contrails.org/reports/v1/reports/v1/v1.html

[81] A. Borella *et al.*, “The importance of an informed choice of CO_{\textrm{2}}-equivalence metrics for contrail avoidance”, *Atmospheric Chemistry and Physics*, vol. 24, no. 16, pp. 9401–9417, Sep. 2024, doi: 10.5194/acp-24-9401-2024.

[82] EUROCONTROL, “EUROCONTROL-CANSO Sustainable Skies Conference: reducing contrails and their non-CO_2 effects on global warming through operational mitigation | EUROCONTROL.” Accessed: Oct. 22, 2026. [Online]. Available: https://www.eurocontrol.int/news/eurocontrol-canso-sustainable-skies-conference-reducing-contrails-and-their-non-co2-effects

[83] Transport & Environment, “How to scale up contrail avoidance in Europe?,” technical report, Jan. 2026. Accessed: Sep. 17, 2026. [Online]. Available: https://www.transportenvironment.org/articles/scale_up_contrail_avoidance

[84] R. Dalmau, G. Jarry, and P. Very, “Contrail-to-Flight Attribution Using Ground Visible Cameras and Flight Surveillance Data.” Accessed: Jul. 01, 2026. [Online]. Available: http://arxiv.org/abs/2510.16891

[85] A. Sarna *et al.*, “Benchmarking and improving algorithms for attributing satellite-observed contrails to flights,” *Atmospheric Measurement Techniques*, vol. 18, no. 14, pp. 3495–3532, Jul. 2025, doi: 10.5194/amt-18-3495-2025.

[86] R. Dalmau, G. Jarry, and P. Very, “Contrails cannot exist without flights: physics-informed contrail detection, tracking, and attribution in ground-based camera video,” *EGUsphere*, pp. 1–38, Aug. 2026, doi: 10.5194/egusphere-2026-4608.

[87] D. P. Nakoneczny, J. P. Itcovitz, D. Vida, L. Busquin, T. H. Abbott, and M. E. J. Stettler, “A flight-first approach for contrail detection using ground-based cameras,” *EGUsphere*, pp. 1–29, Sep. 2026, doi: 10.5194/egusphere-2026-5251.

[88] M. L. Shapiro, “The Cost of Contrail Management.” Accessed: Oct. 16, 2026. [Online]. Available: https://notebook.contrails.org/the-cost-of-contrail-management/

[89] IATA, “Industry Statistics,” technical report, 2024. [Online]. Available: https://www.iata.org/en/iata-repository/pressroom/fact-sheets/industry-statistics/

[90] J. A. Powell, W. Ake, and S. Menendian, “Targeted Universalism: Policy & Practice.” Accessed: Jun. 24, 2026. [Online]. Available: https://escholarship.org/uc/item/9sm8b0q8

[91] “ENR 1.10 FLIGHT PLANNING.” Accessed: Jun. 23, 2026. [Online]. Available: https://www.aurora.nats.co.uk/htmlAIP/Publications/2026-04-16-AIRAC/html/eAIP/EG-ENR-1.10-en-GB.html

aiazero.org

© 2026 University of Cambridge | Aviation Impact Accelerator