FLIGHT SAFETY FOUNDATION

NINE VISIONS, ONE AIRSPACE

Mapping Future Capabilities to Global Expectations for Safety



CONTENTS

1.0 INTRODUCTION **4**

1.1 Terminology 4

1.2 Methodology 6

2.0 WORKS CONSIDERED **7**

2.1 CANSO CATS Overview 8

2.2 Boeing, Skygrid, Wisk Concept of Operations for Future Skies Overview 8

2.3 The U.S. DOT Advanced Air Mobility National Strategy 2025–2036 Overview 9

2.5 The EASA, EUROCONTROL, and SESAR JU Collaborative Digital European Sky – ATM Master Plan 10

2.6 The Flight Safety Foundation IASMS Research Roadmap 11

2.7 AUVSI Blueprint for Autonomy 12

2.8 The FAA HATM ConOps v2.0 12

2.9 EASA AI Roadmap 2.0 13

3.0 A COMPARATIVE ANALYSIS AMONG VISIONS **14**

3.1 Convergence 15

3.2 Divergence 17

4.0 TRANSFORMATIONAL CAPABILITIES **20**

4.1 Five Domains of Focus 20

4.2 Capability Breakdown 20

5.0 MATURITY DEFINED EPOCHS AND FUTURE RESEARCH **25**

5.1 Epoch 1 — Limited (Now–2031) 27

5.2 Epoch 2 — Expanding Autonomy (2031 – 2037) 27

5.3 Epoch 3 — Mature, Seamless Airspace (2037–2045) 28

6.0 FUTURE RESEARCH **28**

1.0 INTRODUCTION

Airspace evolution is constant and continues the process of integrating novel technologies to enable new capabilities that advance safety, security, efficiency, and the economics of aviation services. The latest stages of airspace evolution can be said to be unique in their speed of change, scale of transformation, and scope of impact. Therefore, continuously reassessing the mechanisms and characteristics that map the future of these new capabilities is fundamentally important. While there is a broad variance in our ability to understand when these new capabilities will become operationally available, the scale and scope of the necessary transformations described in various future vision documents are similar across industry, government, academia, and associations. We can describe these transformations with reasonable accuracy and identify and assess the likely safety implications that they introduce. This analysis focused on enabling autonomy in the airspace and therefore removed references to capabilities or goals that explicitly related to non-autonomy functions. This work endeavors to review existing concept of operations (ConOps) and vision documents of future airspace from across stakeholder groups and to draw out their similarities, dissimilarities, and key capabilities to inform industry and government collaboration regarding potential new safety hazards and areas that require a risk assessment.

To inform the future of airspace, it is vital to consider the key players driving change, in addition to those who may not support these transformations. Throughout the research process, industry groups and government leaders highlighted the need for new technologies and new regulatory requirements that would *digitize* the airspace, enable new forms of communication among aircraft, and review airspace access requirements to foster a more information-centric national airspace. In the interest of achieving the level of safety necessary to enable the number and types of aircraft and operations in civil airspace to increase, we seek to understand the technological and regulatory changes driving the long-term evolution of civil airspace.

1.1 Terminology

The following terms are used throughout this analysis. Where definitions are drawn from ASTM F3341/F3341M-24, Standard Terminology for Unmanned Aircraft Systems, the source is noted. Additional terms specific to this analysis are defined to clarify distinctions that are often conflated in policy and technical literature.

Where terms from these sources differ from one another, this analysis uses the definition from ASTM F3341/F3341M-24 as the baseline unless otherwise noted and explicitly identifies where frameworks adopt materially different definitions. The distinction between automatic and autonomy is treated as analytically important throughout this report: Automatic refers to rule-based predetermined execution, while autonomy implies adaptive decision-making in novel situations.

Artificial intelligence (AI) constituent assurance, n — a certification methodology in which an artificial intelligence or machine learning (ML) system is bounded into a discrete, functionally defined unit — the "AI constituent" — for the purpose of airworthiness approval. Rather than certifying the AI system, the AI constituent framework isolates the scope of the AI's safety contribution, defines its operational boundaries and failure conditions, and applies a means of compliance proportional to the criticality of the function performed. The approach resolves the core challenge of certifying a probabilistic, non-deterministic system within a deterministic airworthiness framework by focusing scrutiny on what the AI does and what the consequences of its failure would be, rather than on how it produces its outputs. (EASA AI Roadmap 2.0, 2023)

Artificial intelligence (AI), n — a broad term encompassing computer systems that perform tasks that typically require human intelligence, including learning, reasoning, problem-solving, perception, and decision-making. In aviation contexts, AI is an umbrella concept that includes both automation and ML. AI does not imply autonomy; a system may be AI-enabled while remaining under close human supervision. AI may facilitate autonomy, but autonomy may be achieved without the implementation of an AI-enabled system.

Automated flight rules (AFR), n — a proposed third flight rule paradigm, distinct from visual flight rules (VFR) and instrument flight rules (IFR), in which certified onboard automation manages aircraft separation via onboard systems and third-party

service providers, without dependence on air traffic control for tactical separation. The equivalent European concept is digital flight rules (DFR), see term below. (Boeing, Skygrid, Wisk AFR ConOps, 2025)

Automatic, n — the execution of a predetermined set of procedures or responses by a system without direct human input at the time of execution. Automatic systems follow rules or logic established during design and do not learn or adapt from operational experience. (authors; see also complex function, ASTM F3341/F3341M-24, ASTM F3060-25)

Autonomous, adj — the characteristic of an entity that can, and has the authority to, independently determine a new course of action in the absence of a predefined plan to accomplish goals based on its knowledge and understanding of its operational environment and situation. Having the ability and authority to make decisions independently and self-sufficiently. (ASTM F3060-25)

Beyond visual line of sight (BVLOS), n — an unmanned aircraft system (UAS) operation in which the individual responsible for controlling the flight of the UAS cannot maintain direct unaided visual contact with the aircraft, but the aircraft's location is known through technological means without exceeding the performance capabilities of the command and control (C2) link. (ASTM F3341/F3341M-24, F3442/F3442M)

Command and control (C2) Link, n — safety-critical radio frequency (RF) link(s) between the ground control station (GCS) and the unmanned aircraft (UA). The C2 link is distinct from the data links used for payload or surveillance; its loss triggers defined lost-link contingency procedures. (ASTM F3341/F3341M-24, ASTM F3002)

Complex function, n — a software function or algorithm that may cause a UAS to operate in a manner that is difficult to predict due to compounded implications from factors such as sensor measurement precision, algorithm complexity, environmental variables, multi-core processing, probabilistic algorithms, fuzzy logic, machine learning, genetic algorithms, resource availability, and aircraft system state. These functions are sometimes referred to as "autonomous," "non-deterministic," "artificial intelligence," "adaptive," or "intelligent" algorithms. (ASTM F3341/F3341M-24, ASTM F3269)

Detect and avoid (DAA), n — a subsystem within a UAS providing the situational awareness, alerting, and avoidance functions necessary to maintain safe BVLOS operation of the own ship in the presence of intruders. DAA encompasses both the detect function (temporal and spatial awareness of intruders) and the avoid function (flight guidance to maneuver clear of hazards). (ASTM F3341/F3341M-24, ASTM F3442/F3442M)

Digital flight rules (DFR), n — the European equivalent of automated flight rules (AFR); a proposed operational concept and eventual flight rule paradigm in which unmanned and autonomous aircraft operate under machine-managed separation using digital services, eliminating dependence on traditional air traffic control for tactical deconfliction. DFR is under active development through Project DiFeRent, launched by the European Commission in 2026. (EU Digital European Sky ATM Master Plan, 2024; EASA AI Roadmap 2.0, 2023)

Fail-functional autonomy, n — a design approach that does not rely on a human reversionary mode. Fail-functional autonomy means that the system will automatically continue to operate, perhaps at a reduced capability in the event of a failure without reliance on direct human intervention. Such a design may form the basis of a certification approach for nondeterministic autonomous systems. (AUVSI Blueprint for Autonomy, 2023)

Machine learning (ML), n — computational methods that enable systems to learn without specific programming through the use of statistical, mathematical optimization and other techniques to find patterns in structured or unstructured data sets (that is, training data). The EASA AI Roadmap 2.0 defines a five-level ML taxonomy ranging from ML assistance (decision support with human primacy) through full ML autonomy (complete autonomous decision-making). (EASA AI Roadmap 2.0, 2023,)

Operational suitability, n — the demonstrated fitness of an AI constituent to perform its intended function safely and reliably within a defined operational environment, including after deployment. Operational suitability encompasses post-certification monitoring, performance validation against operational data, and the conditions under which an AI constituent may continue to be used, must be retrained, or must be withdrawn from service. It extends the airworthiness approval beyond

initial certification to account for the capacity of learning AI systems to drift from their validated performance envelope through exposure to operational data not represented in the original training set. (EASA AI Roadmap 2.0, 2023; ASTM F3341/F3341M-24, complex function,)

Remote pilot-in-command (RPIC), n — the person who is directly responsible for and has final authority over the operation of a UAS; who has been designated as RPIC before or during the flight of the UAS; and who holds the appropriate civil aviation authority certificate for the conduct of the flight. (ASTM F3341/F3341M-24, ASTM F3298; 14 CFR Part 107)

Run-time assurance (RTA), n — a system of pedigreed components that implements real-time monitoring, prediction, and fail-safe recovery mechanisms to bound the flight behavior of a non-pedigreed complex function, ensuring the safety of a UAS. The RTA architecture is composed of a safety monitor, a recovery control function (RCF), and an RTA switch. (ASTM F3341/F3341M-24, ASTM F3269)

Strategic deconfliction, n — the arrangement, negotiation, coordination, and prioritization of intended operational volumes, routes, or trajectories to minimize the likelihood of airborne conflicts between operations. (ASTM F3341/F3341M-24, adapted from ICAO UTM Framework)

UAS traffic management (UTM), n — a federated set of services operated under regulatory oversight that supports safe and compliant UAS operations. UTM provides digital services, including flight authorization, strategic deconfliction, conformance monitoring, and constraint management, enabling large-scale BVLOS and urban UAS operations without direct air traffic control (ATC) involvement for each operation. (ASTM F3341/F3341M-24, ASTM F3548)

Unmanned aircraft (UA), n — an aircraft operated without the possibility of direct human intervention from within or on the aircraft. (ASTM F3341/F3341M-24; 14 CFR 107.3)

Unmanned aircraft system (UAS), n — an unmanned aircraft and all required associated elements (including communication links and the components that control the unmanned aircraft) that are required for the safe and efficient operation of the unmanned aircraft in a national airspace system. (ASTM F3341/F3341M-24, ASTM F3178)

Visual line of sight (VLOS), n — an operation in which the pilot or visual observer maintains direct unaided visual contact with the unmanned aircraft throughout the flight to determine its movement relative to intruding aircraft, obstacles, and terrain; observe the airspace for other air traffic or hazards; and determine that the aircraft does not endanger the life or property of another. (ASTM F3341/F3341M-24, ASTM F3178)

1.2 Methodology

To develop a robust and comprehensive review of the existing literature from leading institutions across the globe focused on airspace innovation, the team proposed 11 documents for consideration in this analysis to Flight Safety Foundation's Advanced Aviation Advisory Committee (AAAC). Those 11 initial documents originating from reputable aviation organizations, publicly available, timely, with significant standing in the aviation community, were pared down to nine following team review, discussion with the committee, and analysis of timelines and goals of the proposals. Strategy documents published by the International Civil Aviation Organization (ICAO) were separated from the initial review to be included as a follow-up comparative analysis after identifying the industry and government identified technologies and enabling research necessary for achieving safe and secure autonomy in aviation.

Initially, 57 capabilities were identified in the nine documents included in this analysis, of which 38 were retained after input from the Future of Airspace working group within the AAAC, and internal review by subject matter experts. These 38 capabilities, vetted for accuracy and appropriateness, were analyzed using an AI-based tool built upon the Claude AI platform, using the Sonnet 4.6 Model, for relationships to one another, risks to research prioritization, timeline alignment, alignment across institutional documentation, and support for convergence and divergence analysis. Subject matter experts used this analysis as guidance for further research, for review of any relationships or alignment between the global visions, and to

develop visuals based upon these findings. In all, validation and testing with subject matter experts (SMEs) consisted of over 20 distinct interviews with members of the AAAC and beyond.

The result was 38 capabilities, validated by SMEs and the AAAC, that spanned nine separate but overlapping visions for the future of aviation operations and air traffic management. Dependencies among these 38 capabilities, either explicitly stated or implied, ultimately lead to an initial roadmap. These dependencies rely upon different maturity levels for each capability, and while none can be accomplished or have integrated impact without other mature technologies, some are more closely connected than others. Some technologies can only begin to mature as others enter later states of implementation, and some cannot be realized at all without multiple technologies that enable their success. Among the technological capabilities with dependencies on other, matured capabilities are advanced autonomy for deconfliction, trust in a balanced and safe airspace management framework, high-altitude airspace management that is adequately connected to lower altitudes, and predictive safety systems that drive significant reductions in unsafe conditions.

Maturity definitions were likewise vital to understanding the evolution that will define airspace management over the next 20 years. Many of the experts the Flight Safety Foundation analysis team met to discuss the future reiterated that timelines are less important than technological throughlines – meaning that maturity in the capabilities are what will define each "epoch" in aviation transformation. It is with this in mind that the team developed a maturity map that seeks to define three distinct epochs in the evolution of safe airspace through the lens of increasingly mature technological capabilities as a framework. This paper defines these epochs and their characteristics without specifying a timeframe. Future work will focus on exploring each epoch as a potential future airspace operational scenario to identify unique hazards that may emerge.

2.0 WORKS CONSIDERED

To determine which operational concept and future vision publications to include in this analysis, the following criteria were considered:

- Relevancy of the originating organization to setting an aviation agenda globally;
- Relevancy and timeliness of the vision document to the evolution or implementation of autonomy; and,
- Recommendation by AAAC members.

With these three criteria set, the nine vision documents included in this analysis are:

- **The Complete Air Traffic System Global Council's ConOps for Future Skies**, published in 2025 by the Civil Air Navigation Services Organisation (CANSO);
- The Boeing, Skygrid, Wisk **Concept of Operations for Automated Flight Rules**, published in 2025;
- The U.S. Department of Transportation (DOT) Advanced Air Mobility Interagency Working Group's **Advanced Air Mobility Strategy: A Bold Policy Vision for 2026–2036**, published 2025;
- The U.S. Federal Aviation Administration (FAA) **Roadmap for Artificial Intelligence Safety Assurance**, published 2024;
- The European Union Aviation Safety Agency (EASA), EUROCONTROL, and Single European Sky Air Traffic Management Joint Undertaking (SESAR JU) collaborative **Digital European Sky – ATM Master Plan**, published 2024;
- The Flight Safety Foundation **In-Time Aviation Safety Management System (IASMS) Research Roadmap**, published 2023;
- The Association for Uncrewed Vehicles Systems International (AUVSI) **Blueprint for Autonomy**, published 2023;

- The U.S. Federal Aviation Administration (FAA) Higher Airspace Traffic Management (HATM) **ConOps v2.0**, published 2023; and,
- The European Union Aviation Safety Agency (EASA) **AI Roadmap 2.0**, published 2023.

2.1 CANSO CATS Overview

The CANSO Complete Air Traffic System (CATS) Global Council published its Vision and Roadmap, titled "Collaborating on Future Skies," in 2022. Members of the Council include new entrants like Joby, Indra, Lilium, Wing, Volocopter, Altitude Angel, and Unifly, as well as legacy aviation incumbents like Northrop Grumman, Lockheed Martin, and Airbus. Additionally, the community included regulators, air navigation service providers (ANSPs), international organizations, and nonprofits with interests either in technology adoption or aerospace.

The scope of this roadmap is described as, "The CATS Global Council 2045 Vision imagines global airspace that is safe, fair, intelligent, and interoperable, leveraging revolutionized design, technology and services to power sustainable global mobility and prosperity …" and includes four key operating principles — scalability, interoperability, adaptability, and sustainability. Most importantly, the transition envisioned by the CANSO Vision is supported by systemwide information management (SWIM), which "will be utilizing aircraft with fully connected cockpits enabling large amounts of data to be transferred in real time, transforming service capabilities." [1]

Seven building blocks define the vision's categorization for policy focus. These seven are: services and service delivery; technology and digitization; safety and regulation; environmental sustainability and social impact; airspace design and classification; data and security; and people, organization, and talent. The roadmap follows these seven categories and translates them into 14 long-term goals, which the FSF team has aligned with the Sky for All modernized airspace requirements.

The intent of this roadmap is to "galvanize the industry and to enlist and engage stakeholders in the CATS vision for future skies and the necessary activities to get there."

2.2 Boeing, Skygrid, Wisk Concept of Operations for Future Skies Overview

The Boeing, SkyGrid, Wisk ConOps for Future Skies leverages work and perspectives held by three Boeing-owned companies working on developing technologies to support long-term aviation operations. The focus of this document is to outline a cohesive and collaborative future that also aligns with leading work published by the U.S. National Aeronautics and Space Administration (NASA) and other organizations. Rather than a discussion of "what will aviation look like," this ConOps describes "how aviation will accomplish the task ahead" by 2045.

The evolution from visual flight rules (VFR) and instrument flight rules (IFR) as currently constituted to automated flight fules (AFR) that rely upon "certified onboard automation manage[ing] separation via onboard systems and Third-Party Service Providers …" is fundamentally important to this vision. This parallels a European led project, Project DiFeRent, focused on digital flight rules (DFR) to safely integrate UAS, remotely piloted aircraft systems (RPAS), and autonomous aircraft. Most notably, this proposal is intended to focus on the international harmonization and standardization updates under ICAO and parallel proposals from EASA and FAA, and looks to move forward at that level.

[1] CANSO CATS Pg3

2.3 The U.S. DOT Advanced Air Mobility National Strategy 2025–2036 Overview

With 40 recommendations across six pillars, the DOT AAM National Strategy is a strategic vision for revolutionizing low altitude airspace through industry collaboration and government support. With an emphasis on economic opportunities, the work cites the need for national electric vertical takeoff and landing (eVTOL) demonstrations leveraging existing regulatory frameworks by 2027, to be followed by rapid certification of new vehicle and propulsion types.

The six pillars identified in this strategic document for achieving rapid iteration while maintaining safety are airspace, infrastructure, security, community, workforce, and automation.

2.4 The FAA Roadmap for Artificial Intelligence Safety Assurance

The FAA Roadmap for AI Safety Assurance published in 2024 provides the foundation for authorizing, certifying, or otherwise approving ML and other types of AI models. The roadmap is broken down into seven guiding principles, described in Table 1 below.

Table 1: Seven Guiding Principles for AI Safety Assurance [2]

Work Within the Aviation Ecosystem: Use the existing civil aviation safety requirements and processes while modifying existing methods only when the existing pathways are found to be inadequate for AI characteristics.
Focus on Safety Assurance: Do not sacrifice aviation safety when applying AI to safety and do not focus on the capability gains of the systems.
Do NOT Personify AI [3]: Treat AI as a tool, not a human, by focusing on clear responsibility assignment and avoidance of human-centric language to maintain a clear understanding of the role and limitations of AI in aviation.
Differentiate Learned vs. Learning AI: Both types of AI systems — learned and learning — require different approaches for safety assurance. For learning models, continuous monitoring and runtime assurance are necessary, while learned models can be fully vetted prior to operation.
Use Existing Regulation First: Avoid parallel systems of certification that can fragment the regulatory landscape, and leverage existing policies such as ARP4754, DO-178-C, and DO-254 before developing new approaches of oversight.
Incremental Approach: Mirroring the "crawl, walk, run" iterative process; start with low-risk applications such as forensic or prognostic analytics and predictive maintenance before approving high-risk applications such as air traffic automation or safety of flight AI decision-making.
Leverage Industry Consensus Standards: Avoid developing parallel certification structures across multiple nations by working with standard organizations such as RTCA, SAE, EUROCAE, and ASTM International for national certification.

[2] https://www.faa.gov/media/82891 Section 7

[3] https://www.faa.gov/media/82891 pg 3

2.5 The EASA, EUROCONTROL, and SESAR JU Collaborative Digital European Sky – ATM Master Plan

The 2025 edition of the European Air Traffic Management Master Plan outlines the path forward for ensuring European skies lead in efficiency and environmentally friendly implementation by 2045. The vision is put together by SESAR JU 3, which includes ongoing representation from the European Commission, EUROCONTROL, and 50-plus organizations from academia, standards development organizations (SDOs), regulators, ANSPs, military airspace users, air and ground industry participants, airport leadership, subject matter experts (SMEs), and professional associations.[4] The vision seeks to be both a vision document and an implementation plan for advancing flight in European skies and includes 10 distinct "strategic deployment objectives" intended to be in use by 2035 (Table 2).

Table 2: Strategic Deployment Objectives by 2035 [5]

SDO 1	Alerts for Reduction of Collision Risks on Taxiways and Runways
SDO 2	Optimizing Airport and TMA Environmental Footprint
SDO 3	Dynamic Airspace Configuration
SDO 4	Increased Automation Support
SDO 5	Transformation to Trajectory-Based Operations (TBO)
SDO 6	Virtualization of Operations
SDO 7	Transition Toward Performance of Air-Ground Connectivity (Multilink)
SDO 8	Service-Oriented Delivery Model (Data-Driven and Cloud-Based)
SDO 9	CNS Optimization, Modernization, and Resilience
SDO 10	Enable Innovative Air Mobility (IAM) and Drone Operations

[4] https://www.sesarju.eu/sites/default/files/documents/reports/SESAR%203%20launch%20brochure.pdf pg 5

[5] https://www.sesarju.eu/sites/default/files/documents/reports/European%20ATM%20Master%20Plan%202025.pdf pg 17

2.6 The Flight Safety Foundation IASMS Research Roadmap

The IASMS Research Roadmap is the foundational publication focusing exclusively on NASA's System-Wide Safety (SWS) project within the Aviation Operations Safety Program (AOSP). The research maps the capabilities, technologies, standards, and policies necessary to evolve aviation safety management from a forensic, post-accident discipline into an "in-time" predictive safety assurance system that connects across traditional and emerging aviation operations.

The core elements of the IASMS continue to be organized around the traditional functions of SMS safety assurance and risk management: monitor, assess, mitigate. Within these three main functions of IASMS there are six "swim lanes," identified in Table 3. Four phases are identified within the research to map the evolutionary timelines expected across the six swim lanes; ending in 2045.

Table 3: Research Swim Lanes Within the IASMS Research Roadmap

Safety Data and Resilience Analysis: Core functions for this swim lane include opening the SMS data collection from accidents/incidents to all operations, establishing a safety information exchange model (SIXM), integrating predictive safety management methodologies, and developing predictive systemwide analytics.
Strategic Conflict Management: Core functions of this swim lane include research on scalable deconfliction for BVLOS UAS, evolving from limited UAS traffic management services to fully autonomous strategic conflict management, and coordination of new entrants such as space launch and high-altitude operations.
Tactical Separation Management: Core functions include autonomous self-separation, pairwise separation standards for mixed crewed/uncrewed airspace, and detect-and-avoid (DAA) for uncrewed aircraft systems.
Individual Vehicle Flight Management: Flight management systems needed for 4D trajectory-capable, contingency management for autonomous vehicles, certification for onboard autonomy, and other research for highly automated systems to navigate safely are all included as core functions.
Aviation Weather: Micro-weather for low-altitude UAS, real time weather integration into safety decision loops, and weather-related risk modeling for new entrants are core functions that enable new entrants and BVLOS operations in complex urban environments.
Cross-Cutting Research and Development: Spectrum management, human-automation teaming, cybersecurity for safety critical systems, policy mechanisms, international harmonization, and data governance are all core functions of this swim lane.

This roadmap is unique in that it is exclusively focused on safety research as a prerequisite for more advanced operations such as UAS, UAM, and space-based transit operations. The work highlights specific research gaps that must be closed before the automated, data-driven, multi-entrant airspace can be operated safely on scale.

2.7 AUVSI Blueprint for Autonomy

This blueprint published by AUVSI in 2023 acts as foundational industry guidance for developing a plan to enable autonomous aviation across all levels of flight — localized UAS, BVLOS operations, automated urban flights, and regional air mobility — in a safe and secure way. By leveraging 60-plus organizations across the civilian and military sectors, the blueprint maps out five foundational building blocks and 25 action "bricks" with near-, short-, and medium-term timelines for impact. The document also identifies key policy priorities that act as anchors for enabling new technologies in a productive and cohesive way across the timelines.

The five building blocks and their corresponding 25 action bricks can be found in Table 4.

Table 4: AUVSI Blueprint for Autonomy Foundational Building Blocks

Motivation	Safety and Security, Public Benefit, National Competitiveness, Economic Benefit, Social and Environmental Equity
Technology	Sensing/DAA, Digital ATC, Autonomous Ground Operations, Equipage and Avionics, C2 Communication and Spectrum
Airworthiness	Safety Continuum, Fail-Functional/Run-Time Assurance, AI/ML Certification, Off-Board Systems, Means of Compliance Certification Basis
Operations	Part 135/91/61 Updates for Autonomous Aircraft, Functional Breakdown, RPIC and Multi-Vehicle Supervisor Roles, Lost C2 Link Planning, Continued Operational Safety Definitions
Integration	Automated Self-Separation Assurance/Update to 91.113, Open ATM Service Architecture, Digital Flight Rules (DFR), Legacy Aviation Integration, Local Engagement, Third-Party Service Providers

The AUVSI Blueprint for Autonomy is intended to be the "only industry-consensus document that maps the complete regulatory and certification agenda" needed to enable the types of future visions defined in the CATS ConOps, AFR and DOT AAM Strategy.

2.8 The FAA HATM ConOps v2.0

The focus of the FAA HATM ConOps 2.0 is on outlining solutions to operational challenges of managing air traffic about Flight Level (FL) 600 (about 60,000 ft), which is a domain that is rapidly transitioning from a relatively sparse and passive permissive environment into a commercially congested one. With the testing and envisioned integration of hypersonic transport, high-altitude pseudo-satellites (HAPS), stratospheric UAS, and rapidly increasing rates of commercial space launch and re-entries redefining that airspace, yet without a comprehensive ATC framework, any future concepts for airspace evolution will need to include high-altitude operations.

The HATM is broken into four key concepts to foster short-term and medium-term development of ATC-enabled safety equivalents at FL 600 and above. These four concepts include dynamic corridor management, performance-based access, integrations with national airspace traffic flow management, and scalability for commercial growth. The "seamless, fully integrated global airspace" envisioned by the CATS ConOps is mirrored in the HATM, which holds that future-proof concepts should include high-altitude operations.

2.9 EASA AI Roadmap 2.0

The EASA AI Roadmap version 2.0 provides a structured framework for developing safe certification of AI and ML applications within the aviation community. This roadmap empowers the European regulatory structure to provide AI constituent assurance broken-down into a five-level ML taxonomy intended to set forth a foundation for legislative action on AI deployment. The taxonomy ranges from today's data analytics tools through "fully autonomous airborne" functions. An AI constituent is the approach's fundamental unit of certification and focuses on the functionality of the tool being certified as the determinant factor for the depth of scrutiny required for airworthiness.

In developing the "AI constituent" concept, EASA intends to align this new airworthiness approach with existing systems. Mostly, the focus resolves the challenges of certifying a probabilistic system by applying a deterministic certification framework by isolating and bounding the scope and safety contribution of AI.

Table 5 illustrates the breakdown in the five-level taxonomy.

Table 5: Five-Level Taxonomy for Machine Learning

1. ML Assistance	Decision Support With Human Primacy
2. ML Collaboration	Human-Machine Teaming
3. ML Delegation	Supervised Automation
4. ML Agency	Highly Autonomous Operation
5. Full ML Autonomy	Complete Autonomous Decision-Making

3.0 A COMPARATIVE ANALYSIS AMONG VISIONS

The nine frameworks considered in this analysis were developed independently by different organizations, for different audiences, and in service of different institutional mandates. They are mostly aligned across characteristics in automation as a foundationally important, supported by incremental deployment, with information-sharing as a key tenant of infrastructure, all underscored by an essential role of third-party services. For policymakers, safety researchers, standards development organizations, and civil aviation authorities navigating the rapidly changing landscape of advanced air mobility and autonomous operations, this comparative analysis offers a structured view of where the field agrees, where it may diverge, and where critical gaps remain unaddressed. Understanding both the alignment and the tensions is not an academic exercise: Industry alignment accelerates standards development, reduces regulatory fragmentation, and lowers the cost of certification across borders. Equally, the healthy tensions between frameworks are precisely the pressures that force harder thinking, expose untested assumptions, and drive more resilient technology development than any single organization's roadmap could produce on its own.

This comparative analysis also serves a direct safety risk–assessment purpose. The points of convergence and divergence identified in Section 3.1 through 3.4 could directly inform a preliminary hazard analysis that is intended to be conducted as future work. Where frameworks agree on an architectural requirement — such as trajectory-based operations or machine-to-machine separation assurance — the associated hazards are better scoped and the research priorities more tractable. Where frameworks diverge — as they do on the human role in automated systems, the legal accountability for AI-driven decisions, and the speed of regulatory transition — the discrepancy itself is a signal. The discrepancy identifies where technical capability is outrunning governance, where the safety case remains open, and where preliminary hazard analysis must be conducted with a wider envelope of possible failure modes and system states. In this sense, the map of convergence and divergence is not just descriptive; it is the analytic foundation for the prioritization decisions that follow.

One significant gap cuts across all nine frameworks without exception, and it is the most conspicuous absence in the literature reviewed: operational and information system security. None of the nine source documents directly addresses cybersecurity as a first-order constraint on the architectures they propose. Where cybersecurity appears, it is a supporting consideration or a future research area — not a design prerequisite embedded in the operational concepts themselves. This silence is notable precisely because the architectures where all nine frameworks converge — machine-to-machine separation, systemwide information management, third-party service providers performing safety-critical functions, and remotely operated autonomous vehicles — are architectures whose safety properties are entirely dependent on the integrity of the data and communication layers beneath them. A compromised SWIM network, a spoofed UTM instruction, or a manipulated AI separation decision does not degrade gracefully; it fails in ways that legacy aviation systems, designed around human-in-the-loop cross-checks, were never built to handle. This gap is flagged here and carried forward as a structural limitation of the current policy environment. From an operational perspective, security is likewise undeveloped, and more attention must be paid to the impact that increasing autonomy across an organization has on physical security threats.

3.1 Convergence

1. **Automation as a Central Organizing Principle**

   Each of the frameworks reviewed converges on a similar point, characterized by the emergence of a key technology as a central enabler of how airspace is managed and accessed. The DOT automation pillar, AFR machine-managed separation, HATM's mechanisms for deconfliction, SESAR's Phase D block, EASA's AI Roadmap, and CATS Phase 3's seamless connectedness all treat automation an architectural prerequisite, rather than a feature to be considered.

2. **Incremental Deployment Over "Big Bang" Transitions**

   Though the mechanisms eventually require transformation of the entire airspace across all the source documents, equally prevalent is the expectation that there cannot be one disruptive, rapid-replacement approach. The FAA AI Roadmap, for example, mandates implementation based on low-risk operations similar in nature to the way small UAS operations have given way to larger, more complex, higher risk operations. AUVSI's Blueprint is divided into 25 bricks that are noted as near-term, short-to-medium-term, and medium-term. DOT's LIFT framework (the structure defining DOTs strategic vision for AAM) identifies operational readiness in sequential order separate from time, while the EU Master Plan sets 10 objectives to achieve before 2035. Though they have different timelines, all of these organizations agree that demonstrated safety must come before scale.

3. **Information-Sharing as Critical Element**

   The frameworks included in this study concur that information-sharing is a new and critical element of aviation infrastructure that will drive the necessary transformations. The connective components of this element may be called different names or serve different functions, but whether focused on CATS Phase 1 SWIM data modernization needs; EU Master Plan's focus on cloud-based, data driven service; HATM's higher airspace data integration across the National Airspace System and space operations; or AFR's real-time multi-party data exchange for separation, data connectivity is necessary and ubiquitous, requiring investment and appropriate oversight.

4. **Third-Party Services**

   Though third-party services are not entirely new, all frameworks explored the need for third-party services for separation, navigation, communications, aeronautical information, and intent- and status-sharing in the air and on the ground in a more robust way. DOT, AFR, CAR's, and the EU all rely upon an approved service provider's ecosystem to support the increase in density and complexity, while AUVSI goes a step further to identify a certification schema for achieving and developing that ecosystem.

5. **Safety Research as a Prerequisite for Everything Else**

   Every vision of an increasingly automated airspace outlined in these source documents depends on forms of safety assurance that do not exist today. Developing and refining these safety assurance methods will require continued research that can be applied in real time across a live network. The FSF IASMS Roadmap is the only document that explicitly focuses on the safety research gaps needed to close the safety case upon which the other frameworks rely (SIXM, SWIM, AI Analytical Tools, etc.). Both EASA and FAA make the same argument that the certification layer must be researched, validated, and monitored for increasingly automated system implementation to move forward.

6. Closing the Gap in AI Certification

Artificial Intelligence in aviation is largely defined across two documents: The FAA AI Safety Assurance Roadmap and the EASA AI Roadmap 2.0. These documents have several differences, but philosophically, there is considerable alignment: Phased deployment that begins with low-risk applications, incremental expansion as safety data accumulates, and bilateral agreements on standards will enable certification. FAA's seven principles defined in its framework harmonize well with EASA's five-level ML taxonomy and AI constituent concept with some differences noted in Section 7, below.

7. International Alignment Required

ICAO is tasked with recommending and standardizing practices across the aviation industry and therefore has an immense responsibility for understanding the current technological landscape as well as future innovations that could further enable safety. ICAO's vision and implementation documents are the Global Air Navigation Plan (GANP) and the Global Air Traffic Modernization Operational Concept (GATMOC).

Global Air Navigation Plan (Doc 9750, 7th edition, 2022): The GANP is ICAO's highest-level strategic document to drive the evolution of the global air navigation system. It is structured around the Aviation System Block Upgrade (ASBU) framework which is organized into six-year increments with the most future-focused block upgrade focused on 2031 and beyond.

ICAO's Global Air Traffic Management Operational Concept (GATMOC) (Doc 9854, 2005): Endorsed in 2004, the GATMOC is intended to "guide the planning and implementation of global air traffic management." By providing a description of emerging technologies and targets for the future ATM system, the GATMOC provides a global implementation plan for air traffic solutions.

There is significant alignment between these two documents, within the broader visions for aviation transformation. Trajectory-based operations (TBO), for example, are well-aligned across all visions and ICAO guidance. GANP Block 2 and Block 3 are built around the TBO concept and are likewise defined by a 4D-capable flight management system (FMS), and demand-capacity balancing capability systemwide. The timeline and sequential ordering of the TBO prerequisites align well across ICAO's six-year block timeline.[6]

Surveillance capabilities are also well aligned across the nine source documents and the two ICAO vision-setting guidance documents. Sequentially, the ASBU framework begins with automatic dependent surveillance–broadcast (ADS-B) Out via satellite through future ADS-B In/Out in Block 2 and beyond; this is consistent with the visions outlined by leading organizations.[7]

Open ATM service architectures (including SWIM) map closely with the GANP's generic information management objectives for global integration. The ASBU framework sequences SWIM, other flight information exchanges, and initial flight and flow information (FF-ICE) deliverables. Similarly, RPAS BVLOS operations and UTM/U-Space Service Provision capabilities are aligned closely with the GANP RPAS sequencing. Of the RPAS modules defined in Blocks 1, 2, and 3, national level governments are driving Block 2 integration goals today, and with the publication of Part IV of Annex 6, both ICAO frameworks and the vision documents remain well aligned.[8]

Other components of these international frameworks that are closely aligned include the similarities between Resilient precision and navigation technologies (PNT) and Advanced Precision Approaches enabled by ground based augmentation systems (GBAS) and GANP navigation outlines; aircraft to aircraft (A2A) communication may be more granular in nature than captured in this analysis. For the ASBU communication block, VDL Mode 2, LDACS, and CPDLC define a progressive data link evolution while the capability analysis in this report merges these capabilities.

[6] https://www.faa.gov/air_traffic/technology/adsb/atc/assc
[7] https://www.icao.int/EURNAT/ASBU
[8] https://www.icao.int/aviation-system-block-upgrade-asbu-implementation-monitoring-report

3.2 Divergence

1. **How to Confer Autonomy and When**

   Each framework accepts the eventual handoff of authority to automation but fundamentally disagrees on how quickly to reduce the human role and what residual functions remain for humans to perform. The AFR/DFR-focused elements suggest that no pilot will be necessary in the cockpit for deconfliction as soon as automation can be certified; the human is deliberately designed out of tactical decisions. Meanwhile, the FAA AI publications explicitly require human accountability for "high-consequence decisions" that may include tactical deconfliction. Further, the FAA focuses on expanding AI decision-making from low criticality decisions to increasingly higher criticalities, while a pilot or controller retains legal responsibility throughout. It's important to note that accountability, responsibility, and authority are not to be used synonymously. According to ASTM AC377, accountability refers to the "obligation to answer for an action taken by a responsible entity," authority refers to "the power to give orders and/or make decisions," and responsibility is "the obligation or duty to carry forward an assigned task to its successful conclusion."[9] From a timing perspective, the CATS ConOps sees the elimination or reduction of human strategic oversight occurring beginning in Phase 2 but not complete until Phase 3, or after 2040, when full automation is expected. This evolution represents a much more measured and deliberate transition over two decades. Finally, the EU version draws a future that does not end in full AI accountability, but rather in "human-machine teaming" in Phase D, reinforced explicitly by EU labor law and certification frameworks that reinforce sustained human involvement through at least 2035.

2. **Agility vs. Rigidity**

   Drawing upon the differences between U.S. and European frameworks related to advanced air mobility and U-SPACE certification it's clear that two approaches to enabling future operations are emerging. The U.S. frameworks focus on enabling market-driven experimentation through existing rules and waivers or exceptions while the EU frameworks favor systematic rulemaking and a broader understanding of systemic integration before operations begin at even a small scale. The DOT AAM vision looks to leveraging existing frameworks to enable LIFT Phase L with a stated objective of national competitiveness alongside safety rationale. The EU focuses on promoting an EU regulatory sequence as systematic and multi-year by design to enable U-Space certification as a precursor to larger scale BVLOS operations.

3. **Responsibilities for Controlling Future Airspace**

   Each of the frameworks seems to endorse a third-party service provision for airspace management, but they fundamentally disagree on exactly how much control should be transferred from ANSPs to privately owned and run companies, and on who should ultimately retain final accountability for decision-making. Outlined by the AFR/DFR reports, third-party service providers will provide separation assurance, a service traditionally provided exclusively by sovereign ANSPs. This changes the dynamic of most ANSPs as a standard-setter (or recognizer) and oversight body rather than an operational enabler. Alternatively, the EU U-space vision highlights a more duplicative relationship where third-party service providers operate alongside ANSPs, not taking over sole ownership of separation. The EUROCONTROL network and national ANSPs are more proactively engaged in deconfliction and private services fill gaps that ANSPs assign, auction, or make available otherwise, and full accountability remains with the public institution.

4. **Adaptation vs. Evolution**

   The mechanisms that enable uncrewed aircraft to exist in international airspace continue to be a contentious point globally. Two philosophical approaches remain a point of conversation driving either an "adapt" or "evolve" mindset. On one side are new entrant companies seeking to develop a new class of airspace; this evolution is best exemplified in the AFR/DFR vision referred to as Class X. This Class X would be a purpose-built airspace class for high-density operations enabled through automation. Alternatively, adapting to existing structures is an approach best seen in

[9] https://astmxcellerate.com/wp-content/uploads/2023/10/AC377-Autonomy-in-Aviation-R10.pdf

the EU U-Space vision, where new forms of airspace management are overlaid onto existing airspace classes through dynamic reconfiguration. In this instance, the existing ICAO class structure — without adding new airspace classes — is preserved with digital service layers added to enable increasingly automated operations.

5. Responsibility and Accountability for Incidents Involving Automation

One gap between these frameworks will likely become a major point of contention. The question of who bears responsibility for an incident or for adverse outcomes resulting from decisions made by an autonomous system is yet to be answered, and the frameworks themselves take different stances. The AFR/DFR vision seems to imply that manufacturers and third-party service providers leveraging certified AI models hold some if not all liability in the cases of incidents or accidents while the FAA AI Roadmap identifies the original designer and developer of the AI model as having much of the liability for events caused by AI decision-making. Between the two, the question of who bears responsibility when multiple independently certified autonomous systems, including AI-enabled systems, interact and collectively cause an incident or accident is not sufficiently determined. The EU AI Act of 2024 creates a parallel liability framework for high-risk autonomous decision making orthogonal in some ways to the EASA certification pathway ensuring at least some conflicting accountability for the same systems — a problem that must be addressed prior to widespread implementation and adoption of autonomous decision-making.

When the human has authority and responsibility, it is often clear who is accountable at least at the first line. However, as we delegate authority to automation for decisions and for taking actions, the question of the accountable party becomes more complicated. While a decision-maker like the pilot is accountable for a poor decision, we often perform root cause analysis to determine that the pilot was not alone in making a mistake. The systems the pilot used, the procedures followed, or the training received may have contributed to the incident, which means that other parties can be held accountable.

Accountability is an essential element of a safety management system (SMS), the ICAO-endorsed standard methodology for safety in aviation. SMS is defined as "a systematic approach to managing safety, including the necessary organization structures, accountability, responsibilities, policies, and procedures."[10] Responsibility and accountability are referenced as distinct items in this definition. Responsibility is tied to the execution of a task or activity (and can be delegated) while accountability involves the identification of the person answerable for the outcome of a decision or action (and cannot be delegated).[11]

6. Industry Timelines vs Governance Timelines

A clear discrepancy between the nine frameworks examined in this review can be seen in the timelines associated with the implementation of the various regulatory foundations needed to implement technological changes. The AUVSI Blueprint calls for DFRs, fail-functional autonomy certification, and 1: many operations by 2031 while the ICAO consensus-based standards pace falls far behind that timeline. Similarly, the EU Project DiFeRent (DFR research launched in February 2026) seeks to validate the concept of DFR prior to rulemaking; rulemaking itself takes four to six years in the European Commission. Ultimately, industry timelines endeavor to push the timeline forward with smaller scale implementations and faster go-to-market validation approaches while policymaking led by government and international organizations seems to provide more realistic timelines. (The DOT AAM Roadmap is an exception that focuses on initial short-term wins followed by long-term regulatory development.)

[10] [19] Federal Aviation Administration, Safety Management Systems (SMS), https://www.faa.gov/about/initiatives/sms, Accessed 16 June 2025. 31

[11] International Civil Aviation Organization (ICAO), Annex 19 - Safety Management, 2nd Edition, July 2016.

7. AI Taxonomies for Approval

The FAA and EASA are both looking to certify AI/ML-enabled systems, yet their approach to the tasks is structurally different. The FAA approach defined in the AI constituent assurance framework is principles-based, highlighting seven guiding principles applied flexibly to a wide range of applications. The EASA AI Roadmap 2.0, on the other hand, provides a five-level ML classification system that triggers each level as defined by a set of assurance requirements. Recognizing that having a bilateral certification gap for AI-enabled systems may complicate trans-Atlantic fleet deployment and mutual recognition agreements, the AUVSI Blueprint specifically calls for harmonized AI/ML certification between FAA and EUROCAE.

8. International Misalignment

The most fundamental difference between the ICAO frameworks that provide guidance on the evolution of aviation capability is not technical but architectural. Current airspace operations, established procedures, and existing policies are based on the assumption that the entity ultimately managing airspace deconfliction is a human controller (and/or pilot), even if that person is assisted by automation tools. On the other hand, most, if not all, of the national regulatory and industry frameworks included in this analysis describe futures in which that assumption is no longer true: Third-party service providers, automated self-separation, and multi-aircraft interconnected operations that balance capacity make routing decisions without ATC guidance and are seemingly incompatible with established policies and practices.

AI Constituent Assurance and Operational suitability do not have equivalencies in the GANP, which is why ongoing work has accelerated on these topics. One effort, the Minimum Implementation Path (MIP), is being developed in partnership with the GANP study group to identify and develop practical mechanisms for interoperability in a new digital foundation for future air navigation. The AI certification methodologies, non-deterministic system validation, and AI Constituent concepts are not addressed in ICAO guidance, even while FAA, EASA, and other leading national regulators are beginning to propose, and, sometimes implement, AI certification frameworks.

DFR/AFR and BVLOS rulemaking also have limited, or no, ICAO equivalence. The idea of a third flight rule built specifically for autonomous aircraft that operate outside the traditional VFR or IFR paradigm would require significant changes to ICAO Annex 2 and Annex 6. Though the 2045 visions generally anticipate having a DFR/AFR framework evolving between 2027 and 2033, the process at ICAO (the ICAO Flight Operations panel) would need to begin that process now to support a 2030s deployment; this is not happening.

From a technical perspective, prerequisite capabilities that lead to highly autonomous system deployment require ICAO recognition that does not exist today. Neither the GANP nor GATMOC include modules on fail-function autonomy certification, AI-specific risk management, or runtime assurance for adaptive AI. The current technical safety framework is built upon assumptions stemming from frameworks like DO-178 and ARP4753, and specifically on deterministic certification standards that do not support learning AI systems, which are non-deterministic in nature.

4.0 TRANSFORMATIONAL CAPABILITIES

Considering the variety of opinions, research, priorities, and expectations found within the nine documents considered for this research, Flight Safety Foundation team members, including the SMEs in the AAAC worked to identify, compare, and categorize the capabilities foundational to each vision. Where capabilities were similar, but had different names, the teams combined their definitions. Significant disagreements were noted and examined. This review focused on enabling autonomy in the airspace and therefore removed reference to capabilities or goals for non-autonomy functions. Finally, where a capability falls into a category of supporting an enabling environment for innovation and is not technological in nature (e.g., a regulatory change), it is noted as an enabler.

4.1 Five Domains of Focus

The team grouped the capabilities identified to enable airspace operational transformation that incorporates autonomous aviation operations into five domains. These five domains are as follows:

- **Airspace:** Airspace needed to achieve full integration of increasingly autonomous aircraft operations;
- **Aircraft:** Aircraft, onboard aircraft systems, and associated certifications necessary to achieve full integration of increasingly autonomous aircraft;
- **CNS-I:** Systems that enable communication, navigation, surveillance, and information-sharing;
- **Physical Infrastructure:** Airports, vertiports, ground handling capabilities, charging, and other associated physical elements on the ground necessary to support full integration of advanced air mobility and increasingly autonomous aircraft; and,
- **Operators:** The systems and capabilities associated with organizations responsible for operational control of increasingly autonomous flights.

4.2 Capability Breakdown

Table 6 lists the 38 capabilities grouped into five domains. Capabilities that are implied as prerequisites for each of the 38 capabilities are also discussed While some of these capabilities may begin to enable others downstream, next-order capabilities may not be viable as prerequisites until they are completed, tested, validated, or certified. Due to this complexity, the team created a definition set of each capability as an enabler for others in Table 6. These definitions are aligned with industry's best practice and standards where possible, and while the committee and research team endeavored to be thorough, it is almost certainly the case that as capabilities mature and change, new hazards will emerge and that these will fuel the earliest forms of hazard analysis referenced in the next steps section of this report.

Table 6: The 38 Capabilities Implied by the Nine Documents Analyzed and Their Realization

AIRSPACE

#	Realized When …
1	A UTM/U-space system is operational and legally mandated in the relevant airspace, providing strategic deconfliction, flight authorization, and real-time traffic awareness for UAS operators with machine-to-machine interfaces replacing manual ATC coordination.
2	Automated demand-capacity balancing tools are operational across the ATM network, dynamically adjusting trajectory assignments and slot allocations in response to real-time capacity constraints without requiring manual ATFM intervention for routine imbalances.
3	A regulatory instrument with operational and airspace limitations has been enacted that establishes the certification, operational, and airworthiness requirements for autonomous aircraft operating under digital/automated flight rules in U.S. airspace.
4	A regulatory instrument with operational and airspace limitations has been enacted that establishes the certification, operational, and airworthiness requirements for autonomous aircraft operating under digital/automated flight rules in U.S. airspace.
5	U-space infrastructure services are operational, including dynamic geo-fencing, real-time airspace reservation, network identification, and UAS flight authorization, providing the technical foundation for routine unmanned operations in the defined airspace.
6	A certified (perhaps AI-enabled) conflict detection and resolution system is operational in ATC automation, identifying and resolving flight conflicts without requiring manual controller initiation with validated false-positive and false-negative performance metrics.
7	Certified third-party service providers (and/or U-space service providers) are legally authorized to perform separation assurance functions in defined airspace classes, under a regulatory approval framework defining liability, performance standards, and oversight.
8	A third flight rule category — Digital/Automated Flight Rules — is legally enacted (requiring ICAO Annex 2 amendment), defining the conditions under which pilotless autonomous aircraft may operate in non-segregated airspace without ANSP separation services.
9	A traffic management framework above Flight Level 600 is legally established and operationally active, providing coordinated deconfliction for HAPS platforms, hypersonic vehicles, and commercial space launch/reentry trajectories within a defined governance structure.
10	A recognized authority has established binding pre-launch notification requirements, dynamic 4D hazard volume definitions, and real-time airspace closure and reopening mechanisms that are interoperable with existing air traffic flow management systems and accessible to third-party service providers operating in the affected airspace.

Table 6: The 38 Capabilities Implied by the Nine Documents Analyzed and Their Realization

AIRCRAFT

#	Realized When …
11	A flight management system is certified and operational that can plan, execute, and conform to a precise 4D trajectory (latitude, longitude, altitude, time) with RTA accuracy sufficient for TBO operations across the full flight envelope.
12	All UAS operating in the relevant airspace continuously broadcast a Remote ID signal (identity, position, altitude, operator location) that is machine-readable, network-accessible via a UTM system, and enforceable by aviation authorities.
13	eVTOL aircraft has received full type certification (EASA CS-SC, FAA Special Conditions, or equivalent) and is approved for commercial passenger operations, with a defined simplified operational approval pathway in place.
14	Certified processes are in place for identifying, measuring, and mitigating AI-specific risks (including training-data bias, distribution shift, and adversarial inputs) within the avionics, ATM system development lifecycle, and simplified vehicle operations (SVOs).
15	Avionics and flight-critical systems are certified against cybersecurity threats under an accepted framework (DO-326A/ED-202A or equivalent), with defined security levels, threat models, and post-certification vulnerability monitoring in force.
16	An autonomous aircraft or system has a certified fail-functional architecture — demonstrated to maintain a safe, reduced level of autonomous capability following partial system failure — with no requirement to transfer control to a human pilot.
17	A regulatory framework and technical means of compliance exist for assuring AI-enabled systems that adapt, retrain, or update after initial certification — distinguishing their assurance requirements from those of static ML systems.
18	A complete, harmonized AI Constituent Assurance and Operational Suitability framework is in force (aligned FAA and EASA approaches), enabling AI/ML systems to be approved for flight-critical aviation functions through a defined means of compliance.
19	Autonomous aircraft are certified to maintain their own separation from all other traffic using onboard systems alone, without depending on ATC tactical instructions with the certification basis accepted by the relevant authority as replacing controller-provided separation.

Table 6: The 38 Capabilities Implied by the Nine Documents Analyzed and Their Realization

CNS-I

#	Realized When …
20	ADS-B Out (or functional equivalent) is mandated and operational for all aircraft in the relevant airspace class; ground station network provides continuous cooperative position data to ATC automation with no coverage gaps.
21	High-resolution, certified digital meteorological data (wind, turbulence, icing, precipitation) is delivered directly to aircraft trajectory management systems and ATM automation at the frequency and accuracy required for 4D trajectory conformance.
22	Global aircraft tracking and distress alerting is operational with no more than 1-minute normal tracking intervals and 1-minute autonomous distress tracking, with ICAO Global Aeronautical Distress and Safety System (GADSS) SARPs fully implemented by the operating state.
23	A wide area multilateration system is operational in the relevant airspace, providing continuous aircraft position with accuracy and update rate sufficient to meet the separation assurance requirements of that airspace class.
24	Resilient positioning, navigation, and timing is available using multiple GNSS constellations or other signals plus authenticated backup, with certified countermeasures against jamming and spoofing that have been validated in the operational environment.
25	A certified cybersecurity framework protects the ATM network (SWIM, datalinks, ground systems) against intrusion, spoofing, and denial-of-service, with continuous monitoring and mandatory reporting requirements in force.
26	A ground-based augmentation system is certified and operational at the airport, providing Category I (or higher) precision approach guidance to all equipped aircraft — replacing ILS as the primary precision approach system on that runway.
27	A systemwide information management network is operational and mandatory, providing authenticated, standardized real-time data exchange (trajectory intent, airspace status, weather, surveillance) across all ANSPs, operators, and third-party service providers.
28	A space-based ADS-B relay network provides continuous global coverage including oceanic, polar, and remote regions, making every equipped aircraft visible to ANSPs and UTM systems at all times.
29	A certified digital communications system (LDACS or equivalent) is operational for air-ground instruction exchange, supplemented by authenticated vehicle-to-vehicle and aircraft-to-infrastructure data links enabling cooperative separation and intent sharing without VHF voice.
30	A cross-operator real-time safety data aggregation system (implementing IASMS/SIXM or equivalent) is operational, collecting, normalizing, and distributing safety-relevant operational data from all participants for predictive monitoring and regulatory oversight.

Table 6: The 38 Capabilities Implied by the Nine Documents Analyzed and Their Realization

PHYSICAL INFRASTRUCTURE

#	Realized When …
31	ADS-B Out (or functional equivalent) is mandated and operational for all aircraft in the relevant airspace class; ground station network provides continuous cooperative position data to ATC automation with no coverage gaps.
32	High-resolution, certified digital meteorological data (wind, turbulence, icing, precipitation) is delivered directly to aircraft trajectory management systems and ATM automation at the frequency and accuracy required for 4D trajectory conformance.

OPERATORS

33	An AI-based monitoring system is certified and operational, detecting anomalous aircraft behavior, system degradations, or emerging safety events in real time across a fleet — without requiring human initiation of an alert.
34	Remotely piloted aircraft systems operate routinely BVLOS in non-segregated airspace under a certified operational approval, with a validated safety case covering detect-and-avoid, C2 link integrity, and lost-link contingency
35	UAS operators and the operating state have implemented safety management systems meeting ICAO Annex 19 standards adapted for unmanned operations, including hazard identification, risk assessment, and safety performance monitoring at both operator and state level.
36	One remote pilot (or supervisor) is legally permitted and operationally capable of simultaneously managing multiple autonomous aircraft (m:N ratio), with a certified workload model, supervisory interface, and intervention protocol validated for that ratio.
37	An AI-driven predictive safety management system is operational, aggregating cross-fleet operational data to identify emerging systemic risks and precursor patterns before they produce incidents — with outputs feeding regulatory and operator safety decisions.
38	Real-time threat detection and structured incident response capability is operational across ATM and UAS infrastructure, with defined detection thresholds, escalation procedures, and recovery protocols tested and exercised.

5.0 MATURITY DEFINED EPOCHS AND FUTURE RESEARCH

The epochs defined in this section are intended to label the clustering of capabilities that, after analysis of the dependencies, technology maturity, and timing identified in the source documents, define logical evolutionary steps. These epochs are intended to help describe scenarios that, in future work, can be leveraged for safety risk assessment and the preparation of a preliminary hazard analysis.

The Capability Maturity Table describes these three moments in time (epochs) where the technological capabilities either begin, are being implemented, or have completed initial implementation and are being used in the airspace (M1, M2, M3). Each scenario following this maturity table describes the epoch in depth, and while it is the team's best assessment of the nine vision documents, it is expected that there will be disagreement among stakeholders as to the exact level of maturity at each stage.

In the three epochs discussed, each capability is identified as to what level of maturity it will obtain during that period:

- **M1 – Limited:** Implementation has begun and there is only limited availability of the capability;
- **M2 – Expanding:** Implementation is well underway and there is expanded availability of the capability; and,
- **M3 – Mature:** Initial Implementation has been completed and the capability as matured and is broadly available.

Note the dates used in the table to define each epoch are approximate and are used to communicate approximate expectations. They should not be interpreted as defining specific start and end times.

As a final note, the team felt it necessary to emphasize the need for industry consensus standards and policymaking to enable any of these technological capabilities to scale to meaningful operations. Lessons learned from the development of UAS activities demonstrate a requirement for collaboration and cooperation in the interest of broader digital integrations; an example of this can be seen in the way unmanned traffic management is beginning to positively impact scaled low-altitude complexity.

Table 7: Capability Maturity Table

Capability Maturity Table — 38 Capabilities × 3 Epochs

M1 Limited · M2 Expanding · M3 Mature — all capabilities reach M3 by Epoch 3, completing seamless airspace

Capability	Limited Epoch 1 · Now – 2031	Expanding Epoch 2 · 2031 – 2037	Mature Epoch 3 · 2037 – 2045
AIRSPACE			
UTM / U-Space Provision	M3	M3	M3
Demand-Capacity Balancing Automation	M3	M3	M3
Autonomous Aircraft Rulemaking	M2	M3	M3
Trajectory-Based Operations (4D TBO)	M2	M3	M3
U-Space Infrastructure & Geo-fencing	M2	M3	M3
Digital Conflict Detection and Resolution	M1	M2	M3
Third-Party & U-Space Service Providers	M1	M2	M3
Digital/Automated Flight Rules (DFR/AFR)	M1	M2	M3
High-Altitude Corridors & Traffic Mgmt	M1	M2	M3
Space Launch & Reentry Traffic Integration	M1	M2	M3
AIRCRAFT			
4D-Capable Flight Management Systems	M3	M3	M3
Remote ID & e-Conspicuity	M2	M3	M3
Simplified Vehicle Operations (eVTOL)	M2	M3	M3
Avionics Cybersecurity Certification	M2	M2	M3
AI-Specific Risk Mgmt & Bias Mitigation	M1	M2	M3
Fail-Functional Autonomy & Runtime Cert.	M1	M2	M3
Runtime Assurance for Adaptive AI	M1	M2	M3
AI Constituent Assurance & Op. Suitability	M1	M2	M3
Automated Self-Separation Assurance	M1	M2	M3
CNS-I			
ADS-B Ground Network (Out/In)	M3	M3	M3
Expanded Precision Meteorological Info	M3	M3	M3
GADSS / LADR Global Distress Tracking	M3	M3	M3
Multilateration & WAM Systems	M3	M3	M3
ATM Network Cybersecurity Frameworks	M2	M3	M3
Advanced Precision Approach (GBAS)	M2	M3	M3
Open ATM Service Architecture (SWIM)	M2	M3	M3
Space-Based ADS-B & RPAS Tracking	M2	M3	M3
V2V / A2X / A2A Communications	M2	M3	M3
Resilient PNT	M1	M2	M3
Real-Time Safety Data (IASMS/SIXM)	M1	M2	M3
PHYSICAL INFRASTRUCTURE			
Vertiport Infrastructure Standards	M2	M3	M3
Autonomous Surface Operations	M1	M2	M3
OPERATORS			
RPAS BVLOS Operations	M2	M3	M3
SMS Processes and Procedures	M2	M3	M3
In-Time Anomaly Detection	M1	M2	M3
Multi-Aircraft Operations (m:N Ops)	M1	M2	M3
Predictive Safety Management (AI Analytics)	M1	M2	M3
Threat Detection & Incident Response	M1	M2	M3

M1 — Limited: Initial implementation available for a limited set of flights, geographic areas, or airspace classes.
M2 — Expanding: Utility of the capability is growing and its use is increasing while limitations are being overcome.
M3 — Mature: The capability as currently envisioned is ubiquitously available.

By the end of Epoch 3, every one of the 38 capabilities reaches M3 — the point at which seamless, pilotless airspace integration is achieved.

5.1 Epoch 1 — Limited (Now–2031)

Epoch 1, designated as the *Limited Scenario,* represents the early evolution of ATM and the integration of mixed-mode operations similar to today's ATM design in some of the most advanced operations. The digital infrastructure underpinning airspace operations continues to evolve and mature to enable increasing commercial space launches aided by government coordination. Systems like ADS-B, the global aeronautical distress and safety system (GADSS) distress tracking, multilateration, resilient PNT, and precision meteorology are becoming fully deployed and trusted at scale (M3). Around an emerging digital backbone, a layer of operational capabilities (UTM/U-Space provision, demand-capacity balancing, 4D trajectory-based operations) are expanding into initial use (M2) demonstrated today in the UTM deployment in Dallas, Texas. However, deeper scaled autonomy and the regulatory validation processes, which would require systemwide risk management, fail-functional certification, or runtime assurance for adaptive systems, remain at the foundational stage (M1). Trust in this epoch is still anchored in human oversight and accountability; essentially, automation assists decision-making but does not yet carry widespread authority to act independently.

This epoch is characterized by sporadic and unconnected operations across the airspace that are unable to evolve or expand because coverage is geographically and operationally bounded. U-Space infrastructure and BVLOS operations exist in defined corridors and pilot regions rather than as a default condition of the airspace. Rulemaking for autonomous aircraft is underway but not yet finalized in most jurisdictions, and cross-border recognition of autonomous operations is inconsistent. Simplified vehicle operations are likely certified under early special-conditions pathways rather than harmonized international standards.

The mechanism for deconflicting flight operations within this epoch is predominantly one-operator-to-one-aircraft, with BVLOS permitted by the Part 108 rule in the U.S., within defined corridors, or in some regionally enabled areas that enables limited multi-aircraft operations. In positive control airspace, conflict detection and separation still rely on human controllers and remote pilots-in-command; digital conflict resolution and self-separation do not yet exist in certified form. Those standards and rules are in development. Aircraft can fly 4D-optimized trajectories and communicate via vehicle-to vehicle (V2V) or anything-to-anything (A2A) links, but no aircraft is legally permitted to resolve a conflict or tactically separate itself from traffic without a human in the loop, though strategic conflict detection and resolution is being leveraged for assistance. This is an epoch of assisted, supervised flight.

5.2 Epoch 2 — Expanding Autonomy (2031 – 2037)

The capabilities that were coming online in Epoch 1 (cybersecurity frameworks, GBAS, SWIM, V2V communications) reach full maturity and become taken-for-granted infrastructure (M3). The autonomy stack and the regulatory processes that would enable their operation were merely conceptual before, but— is now actively expanding (M2), meaning certified systems exist and are accumulating operational evidence, though not yet at the scale or confidence level required for unrestricted use. Real-time safety data aggregation begins feeding back into operational decisions without human intervention, marking the first epoch where the system starts to learn from its own autonomous performance.

Within this epoch, the operational footprint broadens substantially. U-Space and demand-capacity balancing are becoming default features of managed airspace rather than pilot programs or sporadic implementations. Third-party and U-Space service providers proliferate, and digital conflict detection is available across multiple regions concurrently. The highest-order capabilities however, such as DFR/AFR as a third routine operating mode, high-altitude corridor management, space launch integration, and full automated self-separation are only beginning to be implemented (M2) and are not yet available at scale.

In this epoch, multi-aircraft supervision (m:N operations) becomes unhindered by regulatory limitations, allowing a single remote operator to legally oversee multiple aircraft, aided by predictive safety analytics and in-time anomaly detection. Digital conflict detection is given the authority to act without human oversight in some situations. Aircraft can be authorized for expanding BVLOS operations across wider geographic areas, but automated self-separation, high-altitude transit, and pilotless-by-default operations remain certification targets rather than realized capabilities. Flight is increasingly automated but still supervised.

5.3 Epoch 3 — Mature, Seamless Airspace (2037–2045)

The final epoch through 2045 sees the 38 capabilities in the roadmap reaching full maturity (M3). AI constituent assurance, fail-functional certification, and runtime assurance for adaptive systems have become a settled, trusted foundation for how aircraft are certified and operated. Depth in this epoch is defined by closing every dependency chain that began years earlier (TBO → digital conflict detection → DFR/AFR; V2V comms → fail-functional certification → self-separation). By the end of this epoch, no capability remains gated on another's immaturity.

Operations dependent on mature capabilities are becoming ubiquitous across the airspace. DFR/AFR exist as a fully recognized third routine operating mode alongside IFR and VFR and are globally supported by international standards and recommended practices. High-altitude corridors and space launch/reentry traffic are integrated into the same operational picture as conventional airspace. At this point, there is no meaningful distinction anymore between "pilot program" and "standard operation" — every capability envisioned in the roadmap is available everywhere it was designed to apply.

At this latest stage, the aircraft systems can plan, fly, detect conflicts, and separate from other traffic entirely on their own authority, with no requirement for a human pilot or controller to routinely intervene. Multi-aircraft supervision scales without the practical ceiling that existed in Epoch 2. Flight from the surface through high-altitude corridors into space launch and reentry integration are all governed under a harmonized DFR/AFR framework. This is the epoch in which "seamless airspace" becomes the operating baseline: pilotless, self-separating, fully autonomous flight is the default case.

6.0 FUTURE RESEARCH

With maturity scenarios developed based upon the aggregate evolution and intersection of these 38 capabilities for the future of airspace, the Flight Safety Foundation team proposes to review and explore the unique safety hazards that emerge within each of the three epochs through a structured scenario and vignette-based analytical approach, culminating in a preliminary hazard analysis for each epoch.

This paper establishes the epochs, the capabilities, and the dependency structure among them but intentionally stops short of identifying the specific hazards those configurations introduce. Each epoch represents not just a different set of available technologies but also fundamentally different distribution of authority, accountability, and failure modes across the system. The hazards most likely to cause accidents in Epoch 1, where human oversight remains the primary safety backstop and autonomy is bounded to defined corridors and pilot programs, are categorically different from those that we expect will become prominent in Epoch 2, where the authority to act begins transferring to certified autonomous systems operating across broader airspace, and from those that characterize Epoch 3, where the default condition is pilotless, self-separating flight governed by machine-to-machine authority. A single hazard analysis conducted at any one of these states would not generalize to the others.

The proposed approach draws directly on the scenario structure established in Section 5.0. For each epoch, the research team would develop one or more operational vignettes with narrative descriptions of a specific, representative flight operation that is plausible and characteristic of that epoch's capability maturity level. A vignette for Epoch 1 might describe a BVLOS cargo drone operating in a UTM-managed corridor alongside general aviation traffic in Class E airspace, under Part 108 authorization, with an RPIC supervising a single aircraft. A vignette for Epoch 2 might describe a third-party service provider managing separation for a mixed fleet of eVTOL air taxis and autonomous cargo aircraft across an urban Class B environment, with real-time safety data aggregation feeding automated intervention decisions. A vignette for Epoch 3 might describe a fully autonomous aircraft navigating trans-Atlantic airspace under DFR/AFR, coordinating with space launch traffic through a high-altitude traffic management framework, with no human pilot or controller in the operational loop. The details of each vignette would have to be specific enough to support systematic hazard identification by naming the actors, the data flows, the authority relationships, and the failure modes that are architecturally possible within that epoch's capability configuration,

which will require the team to make a broad set of assumptions for the vignettes set in the far future.

Against each vignette, the team would conduct a preliminary hazard analysis (PHA) following established safety engineering practice, including engagement with SMEs in the aviation community. The PHA would identify credible hazard conditions and assess each against the five accident categories particularly important within the ICAO safety model: runway incursion, runway excursion, midair collision, loss of control–in flight, and controlled flight into terrain. For each hazard, the analysis would document the initiating condition, the causal pathway, the barriers present in that epoch, and whether existing standards, regulations, or operational procedures provide a credible mitigation or whether the hazard represents an unaddressed gap. The capability dependency analysis developed as part of this research will provide a structured foundation for that work. The 38 capabilities identified in this report, their prerequisite relationships, and their domain assignments directly inform which system failures are architecturally possible at each maturity level, and which barriers are or are not in place.

Two findings from the current analysis should structure the prioritization of the PHA effort. First, the comparative analysis in Section 3.0 established that all nine source frameworks are silent on security. This means that for any vignette, cybersecurity threats such as spoofed remote ID, manipulated UTM instructions, compromised SWIM data, or adversarial inputs to AI separation systems must be treated as an unmitigated hazard class unless explicitly addressed. Second, the capability risk analysis identified a cluster of under-researched capabilities whose under-prioritization would not degrade safety gradually but would remove entire categories of barrier simultaneously. The vignette design should deliberately stress-test these capabilities by constructing scenarios in which the capabilities are partially or conditionally available, in order to surface the consequences of development that lag in the research priorities that are most at risk.

The output of this work would be a set of epoch-specific preliminary hazard analyses suitable for use by safety researchers, standards development organizations, and civil aviation authorities as a structured baseline for standards and rulemaking prioritization. Flight Safety Foundation's In-time Aviation Safety Management System research roadmap provides both the methodological framework and the institutional context within which these analyses would sit, connecting the hazard identification work to the broader goal of predictive, data-driven safety management across an increasingly autonomous national and global airspace system.