UK Civil Aviation Authority

The CAA's Response to Emerging AI-Enabled Automation

Part B: Strategy for Using AI in the CAA

CAP3064B

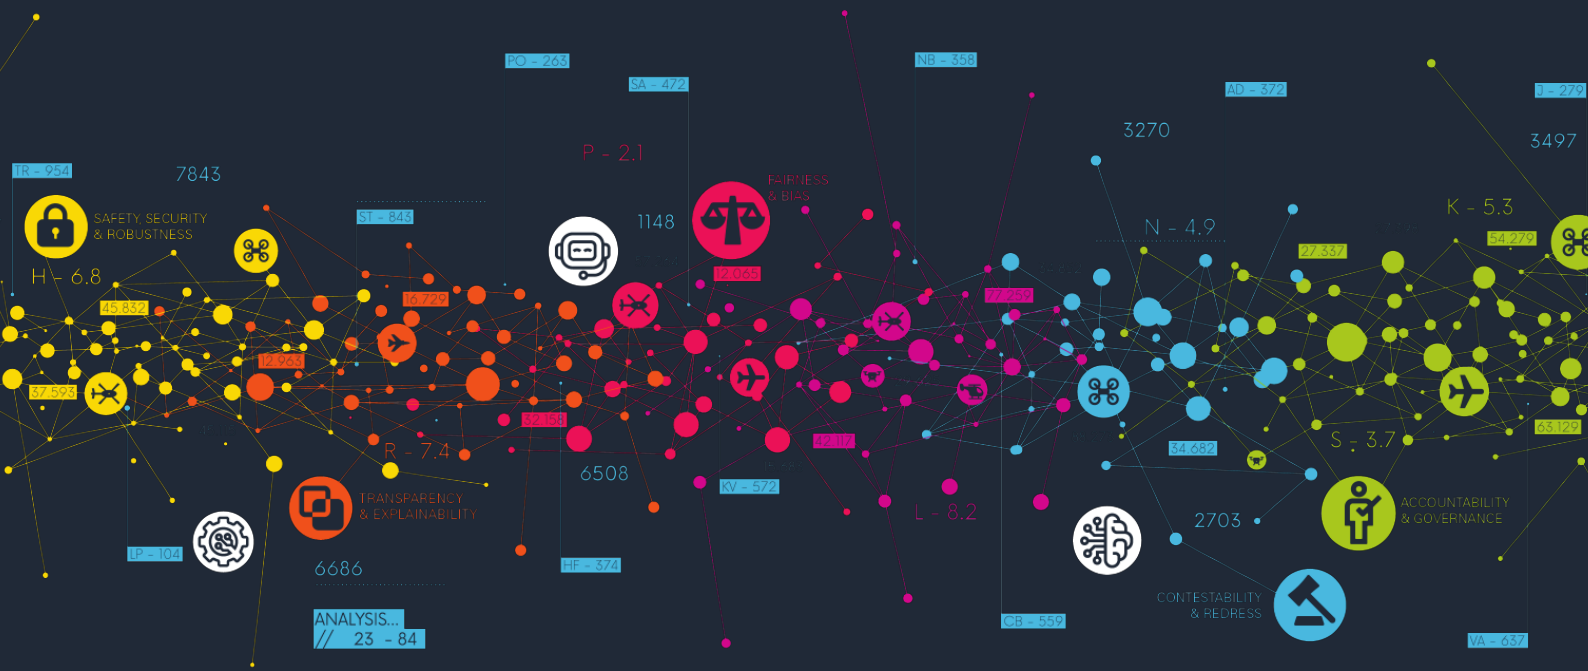

PROTECTING PEOPLE, ENABLING AEROSPACE

CAA Ai Innovation Strategy

v1.3.1 Nov 2024

Published by the Civil Aviation Authority 2024

Civil Aviation Authority
Aviation House
Beehive Ring Road
Crawley
West Sussex
RH6 0YR

First published November 2024

Enquiries regarding the content of this publication should be addressed to: strategyforai@caa.co.uk

The latest version of this document is available in electronic format at: www.caa.co.uk/AI

Executive Summary

This document outlines the UK Civil Aviation Authority's (CAA) strategy for adopting and managing Artificial Intelligence (AI) as a mechanism for enhancing our organisation to deliver the CAA's strategy. Recognising the rapidly evolving landscape of AI and its potential to enhance operational effectiveness, the CAA aims to harness these technologies while mitigating associated risks.

The strategy emphasises a human-first approach, prioritising the enhancement of our colleagues' performance with AI-based automation. It will provide colleagues with guidelines and support for the safe and appropriate use of AI tools, addressing concerns such as data privacy, information security, and the need for human oversight in decision-making processes.

The strategy is founded upon the CAA's AI Framework and its five core principles, which will be integrated into existing CAA governance structures. This framework will guide the development, procurement, and implementation of AI-enabled tools across the organisation. The strategy also builds on our plans for establishing a Community of Practice that will provide expertise and support upskilling across any AI-based initiatives. These are described in our response to emerging AI-enabled automation, within CAP3064, and it is recommended that our strategy for using AI in the CAA is read in conjunction with this.

Our initial focus is to embed this AI strategy throughout the CAA in line with our regulatory principles and organisational values, including within the CAA Strategy, IT Strategy, Customer Strategy, and People Strategy. This alignment ensures that AI adoption supports broader organisational goals while maintaining the CAA's commitment to safety, security, and regulatory excellence. We also introduce the basis for a risk assessment framework specifically designed for AI initiatives, grounded in the AI Principles. This framework will be used to evaluate and mitigate risks associated with AI adoption.

We have set out a phased approach to implementation, starting with establishing the AI Risk Framework, updating policies, and developing training materials in the short term. Flagship demonstrators and proofs of concept will support our experimentation and help colleagues and customer to explore the benefits that the technology has to offer. Medium-term goals include developing comprehensive policies for AI governance and delivering specific policy updates across various business functions. Using illustrative examples as an appendix to the strategy we start to identify potential applications of AI in regulatory, business, and individual colleague contexts.

In summary, this strategy for using AI within the CAA presents a comprehensive and thoughtful approach to AI adoption within the CAA. It balances the potential benefits with necessary precautions, setting a clear path for the organisation to leverage AI

technologies while maintaining its regulatory responsibilities and commitment to safety in the aerospace sector.

Chapter 1

The Strategic Challenge

1.1 Artificial Intelligence (AI) is changing quickly, bringing new opportunities all the time. The UK Civil Aviation Authority (CAA) has a lot of data, and we could use it better to gain more useful insights.

1.2 By taking advantage of the opportunities presented by AI technologies we have the potential to improve the way that the CAA delivers its mission effectively and efficiently. We should not ignore this opportunity to increase the impact the CAA can have on the safety, security, airspace, and sustainability of aerospace for consumers and the public, as well as consumer protection. However, uncontrolled adoption of some of those new capabilities would present a risk to the integrity of the work of the CAA and a risk to our reputation. Equally, choosing not to proactively explore options for the adoption of AI into the CAA may also result in new risks.

1.3 We will define an approach to the selective adoption of new AI technologies, ensuring safe and responsible use, to maximise our operational effectiveness while providing advice, guidance, governance, and constraints to colleagues who use them.

1.4 It should be noted that while AI has recently seen an accelerated increase in interest and investment, many of the challenges it presents are not new. We have both a corporate and societal responsibility to ensure we use any new technology appropriately, effectively, and securely with good oversight and governance. Our challenge is to determine how we most effectively apply those principles to the use of AI by and within the CAA.

Our Vision for Using AI in the CAA

1.5 The CAA's proactive adoption of artificial intelligence will enable colleagues, customers, and consumers to work, collaborate, learn, and operate more effectively.

Chapter 2

Empowering Safe & Responsible AI Adoption

2.1 The CAA's use of AI includes application of the technology for colleagues' everyday use, for our regulatory duties and activities, and for our business functions. In all cases, we will take a human-first approach to ensure that the focus is on improving human effectiveness rather than replacing humans with technology.

2.2 In support of this, we will apply the CAA's AI Principles (available at www.caa.co.uk/CAP2970) to all uses of AI in the CAA and have developed guidelines for colleagues to aid the safe use of AI. This will ensure that we continue to comply with existing legal requirements, and that tools and services used by both colleagues and customers are safe, secure, and trustworthy.

AI Opportunities

2.3 Through engagement with teams across the business and supported by research into the latest applications of AI in other organisations, we have identified several opportunities that provide the CAA with various benefits. To help us communicate, organise, and oversee these we will group use cases into 5 categories:

Modelling

2.4 AI has the power to review large volumes of data and to generate scenarios by creating mathematical or statistical representations of real-world systems – models. The models aim to capture the relationships between various factors, and to predict how they might evolve over time. This allows the model to synthesise possible future scenarios. Applications could include horizon scanning, corporate risks, regulatory risks, and economic regulation.

Creating content

2.5 The ability of a Machine Learning system to find patterns and connections within data also enables it to then generate synthetic data and content that resembles the original source. These tools can generate text, images, videos, music, and audio of a quality equivalent to that which can be created by humans. Applications are extensive, but include voiceovers for videos, new material for training courses, job descriptions, adverts, and interview questions.

Analysing Data

2.6 The CAA holds and is responsible for a collection of data which is not available elsewhere. However, resource constraints and prioritised working mean that

potential remains to more fully extract the possible knowledge and insight from this data such that it can provide additional value to our role in protecting people and enabling aerospace. AI-based tools can interrogate large datasets to extract themes and statistics, recognise patterns, or predict future trends. They go beyond legacy numerical analysis tools such as Excel by also analysing text data from sources like surveys and reports. Potential applications include rapid quality assurance checks on applications, reviewing consultation and survey responses, categorising qualitative safety reports, or sifting job applications.

Managing Knowledge

2.7 Corporate and regulatory knowledge is spread throughout the CAA. We use tools like intranet search and signposting to share and access this knowledge, but the extent of this information can become overwhelming. Using AI, the CAA could maximise its effective use of corporate and regulatory knowledge for the benefits of customers, consumers, and colleagues. With the CAA's recent growth in regulatory responsibilities and therefore the size of the organisation, AI could aid with collaboration, identifying linkages and making connections across the 'system' of the CAA. Opportunities include corporate services chatbot providing quick access to corporate knowledge and policies, management of our regulatory libraries, developing personalised training programmes, and supporting risk identification, assessment, and management.

Making Decisions

2.8 At the CAA we strive to make data-driven decisions. There are opportunities to use AI to enhance our decision making, and to support or automate decision-making processes within the CAA, primarily for low-risk applications. This category also serves as a placeholder for future developments in AI decision-making capabilities. Potential applications include automating routine administrative decisions, prioritising tasks based on urgency and importance, and providing recommendations for resource allocation. This category also covers exploration of the risks associated with AI decision-making and allows for the possibility of expanding into more complex decision-making applications as technology, risk mitigations, and organisational culture evolve. The CAA will carefully monitor and assess the potential for AI to support decision-making in more critical areas, always prioritising safety, security, legal, and regulatory integrity.

2.9 For each of the categories, use cases can apply to regulatory, business, or everyday functions.

Using AI Responsibly

2.10 The CAA's AI Principles (see *CAP2970 "Building Trust in AI"*) are an important guide for all colleagues who are using AI for every-day, business, or regulatory tasks.

2.11 CAA colleagues are encouraged to make use of the opportunities presented by mainstream consumer AI applications, whilst being mindful of the need for care with respect to the limitations and risks associated with them.

2.12 By consumer AI applications we mean those services which employ AI to offer enhanced functionality without the need for specialist knowledge in the user, and without the need for bespoke development or AI model training by the CAA. Examples include ChatGPT, Microsoft Copilot, or Perplexity.

2.13 Considering the benefits and limitations of current consumer tools, we believe that further general guidance is necessary for colleagues until such a time that detailed risk assessments can be carried out on specific tools (as proposed in Chapter 4):

i) Colleagues must not use AI tools for any purpose or activity which would otherwise be considered inappropriate in a work context. To reflect this and to provide guidance on what is considered inappropriate, we will update the CAA Code of Conduct, as well as our Diversity and Inclusion and Wellbeing guidance materials.

ii) Colleagues must not share confidential, sensitive, or personal information that is not appropriate for the public domain with generative AI tools unless they are provided in a secure environment. This is because many tools use information given by users to further enhance their own knowledge and functionality. Some tools may state that user and company data is protected, but it is not always clear whether information fed into these systems is genuinely secure. Colleagues should assume data will not be handled securely in such applications unless the CAA has specifically approved it for sensitive data.

iii) Colleagues must not rely solely on the information, data or calculations derived from AI tools. Information provided by AI tools may be based on the work of others or be the result of unreliable data or calculations. It may therefore be of unknown accuracy and provenance. As with other research sources and analysis tools, it is important to validate the source of information being used, clearly establish the quality of data underlying reports, and to be transparent about the source of any information which may be reproduced elsewhere.

2.14 Guidance, training, communication and support for our colleagues on the responsible use of AI will be a key part of delivering this strategy through the AI

Portfolio. Mainstream consumer AI tools are on a trajectory to evolve rapidly over the coming years, so the above guidance will be regularly reviewed and updated.

Chapter 3

Integrating AI into the CAA

Applying the AI Framework

3.1 It is widely reported that AI has the potential to shift how businesses operate by introducing automation that is more powerful, more accessible, and more integrated. However, as a regulator, the CAA has responsibilities that impact on the safety and security of the public and we must ensure that any use of AI is not detrimental to the fulfilment of those responsibilities.

3.2 The AI Framework introduced in the CAP3064 "The CAA's Response to Emerging AI-Enabled Automation" will enable AI to be adopted, trusted, and used by the CAA and our customers. We must hold ourselves to the same standards we expect of industry when adopting and using AI tools within the organisation, particularly when using those tools to process information relating to regulatory activity and sensitive data.

3.3 Policies and procedures will be written and amended as appropriate to reflect the AI Framework. This will embed the aspirational vision for appropriate adoption of AI across the CAA, alongside the necessary controls to mitigate any risks.

3.4 The CAA's AI principles were first published in CAP2970 in February 2024. It is now important to embed these into colleague training and procedures, so that they become a standardised framework for dealing with AI. We will therefore be exploring how existing training courses within the CAA can be updated to include the AI Framework, from awareness and appreciation of the basic elements for all colleagues, through to more advanced training opportunities for experts. This will require a training needs analysis across all areas of the business to identify the appropriate strategy.

Impact on the CAA as a National Aviation Authority

3.5 As the UK's National Aviation Authority (NAA), we are routinely graded against the International Civil Aviation Organisation's (ICAO) 8 Critical Elements which are necessary for effective safety and security oversight systems and supporting the maturation of our State Safety Programme.

3.6 The CAA's use of AI has potential to impact our ways of working and compliance with ICAO Critical Elements, and so we will ensure that any adoption of AI technology within the CAA is assessed for its impact on these metrics.

3.7 More information is provided in CAP3064 "The CAA's Response to Emerging AI-Enabled Automation".

Integrating with CAA Strategies

3.8 Recognising that AI and the resultant automation are enablers to the delivery of CAA goals, rather than ends in themselves, it is important that we integrate and harmonise the AI Framework with other CAA strategies as appropriate.

3.9 As we progress with our AI initiatives, this strategy will evolve to reflect the changing landscape of AI in aerospace and the developing needs of our organisation.

The CAA Strategy

3.10 The CAA Strategy has 5 strategic themes that include protecting consumers and the public; enabling aviation and aerospace to innovate and grow; developing relationships to improve standards globally; supporting aerospace to improve environmental sustainability; and enhancing our organisation.

3.11 AI technologies have the potential to support each of these strategic themes, but there are particular opportunities with respect to protecting consumers and the public and enhancing our organisation. Any technology which can improve the effectiveness of the CAA in carrying out its role can enhance our organisation and, therefore, improve the degree to which we can protect consumers and the public.

3.12 Appropriate exploitation of opportunities to utilise AI technology safely and ethically is therefore important to ensure we maximise the extent to which we realise our strategic goals.

IT strategy

3.13 The CAA's IT Strategy describes several key themes for the ongoing improvement of our technology estate and effective use of it. The use of AI may affect each of these in ways that we may not yet fully understand. It is therefore important that the IT strategy continues to monitor technological developments in AI and incorporates the AI Framework to ensure that trust is maintained. The CAA's IT Strategy is reviewed and updated at least annually to take account of changing context and opportunity, and the impact of new AI technologies will be part of that review as appropriate.

3.14 The following areas of the IT strategy are of particular interest and relevance to the use of AI and will subsequently be explored further to identify any potential actions.

"Supporting everyday efficiency"

3.15 This theme ensures that colleagues have the right tools and knowledge to safely make the most of the technology available.

3.16 There are expected to be substantial benefits for both the CAA and our customers from adopting AI-based tools. The introduction of AI tools that are readily available to colleagues, such as Microsoft Copilot and ChatGPT, can enhance colleague's ability to work effectively when used in accordance with our guidance. Examples include drafting documents or presentations for specific audiences, analysing qualitative and quantitative consultation feedback, or generating engaging media content to improve stakeholder interest. Further examples are listed in Appendix B.

3.17 Where it is judged to be beneficial, the CAA may also choose to develop bespoke tools using AI technology, as discussed in paragraph 3.30 below.

3.18 Where AI-based tools are adopted by the business, we anticipate that the AI Centre of Excellence will map, monitor, and report these benefits as they arise.

"Enabling the Customer Strategy"

3.19 The CAA delivers many services to the sector and the public, through regulatory applications, advisory services, the ATOL scheme, and more. The CAA's Customer Strategy embodies the ambitions and commitment of the CAA to providing an excellent customer experience that is straightforward, joined up and personal.

3.20 There is an ambition within this to utilise technology that helps us to meet the customer experience goals, but similarly there is recognition of the benefits in retaining human interaction in the services we deliver. There is an ever-growing number of software tools that utilise AI which are designed to enhance customer experience. We will work to ensure that our response to emerging AI-enabled automation, this strategy for using AI, and our AI framework are reflected within the Customer Strategy, and that any introduction of AI-based services is carried out in accordance with the 5 AI Principles and the guidelines set out in this document.

"Enabling the Data Strategy"

3.21 Data is fundamental to the effective use of AI-enabled tools. The Data Strategy is driven by its own core principles for data collection, processing and analysis that align well with the AI Principles. In addition, the CAA's Information Governance Framework clearly describes the relevant roles and responsibilities, legislation, and policies.

3.22 There is therefore a common goal between these 2 strategies regarding the quality and accuracy of data. We will, therefore, work to incorporate this strategy and the AI framework into the Data Strategy to reach this common goal.

"Security Maturity"

3.23 The security maturity objective has enabled the CAA to improve the mitigation of the risks that our IT architecture is subject to. Introducing AI-based tools,

whether bespoke or not, has the potential to broaden existing risks, introduce new risks or reduce those risks.

3.24 Any AI-based tools will be required to adhere to CAA IT security and access control policies alongside appropriate information security legislation. We will also work proactively to evaluate new and existing risks and update or develop appropriate mitigations.

People Strategy

3.25 Our goal is to create a dynamic, AI-ready workforce that can drive innovation while maintaining our commitment to safety and regulatory excellence. As we embark on our AI journey, it is crucial that we develop comprehensive elements for AI within the CAA's People Strategy that address the challenges and opportunities across the organisation. These will focus on three key areas: Building AI, Managing AI, and Adopting AI. Our aspirations are outlined below which will be further explored within the AI Portfolio:

a) Building AI capabilities to establish a strong foundation for AI within the CAA:

- Implementing workforce planning to identify and address skills gaps.
- Conducting skills audits to assess our current AI readiness.
- Identifying and investing in education routes, including industry placements, apprenticeships, and experienced hires.

b) Managing AI responsibly ensuring the responsible use of AI is paramount:

- Reviewing and updating our HR policies to reflect the risks associated with AI.
- Ensuring our learning and development programmes are current, with a particular focus on mandated training.

c) Adopting AI in all areas of the organisation to foster familiarisation and competence:

- Promote AI literacy across all levels of the organisation.
- Integrate AI considerations into our competency and technical frameworks.
- Incorporate AI-related objectives into our Performance Development Conversations (PDC).
- Include AI aptitude assessments in our recruitment processes, such as interviews and psychometric tests.

d) Integrating AI within our People Strategy will ensure that the CAA is well-equipped to build, manage, and adopt AI skills and capabilities effectively, supporting our regulation of AI in aerospace, and our use of AI within the CAA. By focusing on these three key areas, we aim to create a workforce that is not only capable of implementing and using AI solutions but also adept at navigating the ethical and practical challenges they present.

Integrating with the CAA Consumer Strategy

3.26 The CAA's Consumer Strategy outlines our vision for a competitive aerospace market where consumers have access to choice, are well-informed, and have confidence in businesses meeting their obligations. As we develop and implement AI technologies, it is crucial that we align our AI strategy with these consumer-focused goals. This alignment will ensure that our use of AI enhances rather than undermines consumer protection and empowerment.

3.27 Potential crossovers between the AI and Consumer strategies include using AI:

- To improve accessibility and support for consumers in vulnerable circumstances.
- To enhance compliance monitoring and enforcement of consumer protection legislation.
- To empower consumers with better information and decision-making support.
- In economic regulation to further consumer interests in competition and growth.

3.28 To effectively integrate these two strategies, we propose the following initial steps:

i) Conduct a comprehensive comparison and analysis of the 5 AI Principles against the Consumer Principles to assess alignment and extract any additional considerations.

ii) Develop a matrix analysis comparing potential **AI use cases** against the output of the AI and Consumer Principles analysis. This will highlight areas of interest and inform the consumer aspects of our AI Portfolio.

iii) Review the Consumer Strategy's priority focus areas to identify specific opportunities for AI application, such as using machine learning to improve detection of non-compliance with consumer protection legislation or developing AI-powered tools to assist consumers with special assistance needs.

iv) Continue to draw on the extensive expertise of the CAA's Consumer Panel to gather input on the potential benefits and risks of AI implementation

from a consumer perspective, ensuring that consumer interests are central to our AI adoption plans.

v) Develop guidelines for AI-enabled consumer-facing services across the CAA, ensuring they adhere to both the AI Principles and Consumer Principles, with a particular focus on transparency, fairness, and accessibility.

vi) Establish a cross-functional working group with representatives from both the AI and Consumer teams to oversee the integration of these strategies and ensure ongoing alignment as both areas evolve, utilising existing tools such as our Vulnerability Toolkit to guide our focus.

3.29 By taking these steps, we aim to create a synergistic approach that leverages AI to enhance consumer protection and empowerment while ensuring that our AI implementation remains firmly grounded in our commitment to serving the best interests of aerospace consumers.

Developing AI tools for CAA use

3.30 In contrast to every-day consumer AI applications, opportunities also exist to develop new bespoke tools using AI technologies to solve existing business problems.

3.31 To ensure that we learn as an organisation during this early period of adoption and evolution, and to maintain a consistent approach to the responsible use of AI, our Information Services Department (ISD) will retain overall responsibility for any AI-based tool development. In most cases these will be enabled by Proofs of Concept (PoC), allowing us to build, test, and learn through an iterative approach. Governance proposals are described in Appendix C.

"Citizen Development" of AI Tools

3.32 To enable colleagues to identify opportunities within their own areas of expertise, we have our "Citizen Development" programme which is positioned between the simple use of consumer technology by business users and traditional software engineering of bespoke tools. Citizen Development includes the use of end-user technology to develop relatively complex business solutions which go beyond the basic consumer tooling and, if they are considered business critical, require suitable documentation, backup, and security measures. Tools such as Microsoft Access, Excel, SharePoint, and Power Automate are examples which can be used in this way.

3.33 AI tools and technologies will also be able to be used in this way by colleagues to develop valuable business solutions. Governance of these types of solution will be incorporated into the existing Citizen Development processes ensuring that new tools are identified and logged, with any risks managed as appropriate.

Procuring AI-Enabled Tools

3.34 There is a growing prevalence of AI tools on the commercial market. While the CAA takes a mature approach to the selection and implementation of information services across the CAA, the nature of AI solutions may introduce new risks and considerations which may demand an additional layer of analysis.

3.35 Benefits and procurement frameworks for AI solutions are available. To avoid unnecessary duplication, we will test these against potential applications to assess the suitability and relevance for CAA applications, before determining the frameworks that the CAA will ultimately use.

3.36 When procuring new software tools that include an AI element, we will apply the 5 AI Principles in their entirety. To this end the “Non-Functional Requirements” used as part of the specification for any technology procurement exercise will be updated to ensure suppliers explain how those principles will be met by any proposed solution.

3.37 As with all procurement, the starting point must always be the business requirement and not the available solution. Non-AI solutions may be the better fit to satisfy the requirement and it’s important not to become fixated on adoption of a particular tool or technology.

3.38 The CAA often partners with external organisations to support delivery of new digital tools and processes. We will ensure that any supplier or partner complies with the CAA’s policies regarding AI and will review and update our procurement and supplier management frameworks to align to these guidelines and the 5 AI Principles.

Managing Part B with existing governance

3.39 With the AI Framework (see CAP3064 “The CAA’s Response to Emerging AI-Enabled Automation”) as the foundation the use, management, and governance of AI in the CAA will be integrated with existing business governance rather than creating whole new structures unnecessarily. This will enable AI to be incorporated into those existing control structures. We will continue to review the effectiveness of this approach and may make changes where necessary.

3.40 We expect that the CAA will primarily be a consumer of AI tools, rather than a developer of them. However, where a requirement exists for a bespoke solution, we will aim to develop and retain in-house skills, supported by procured expertise as required.

Chapter 4

Managing AI Risks

4.1 Although the standard CAA risk framework will be used to identify, track and share any organisational risks that may arise from the use of AI technology, specific risk assessments will be undertaken as part of the evaluation of opportunities to use AI technologies.

4.2 Because this area of technology is volatile and fast-moving new risks may emerge quickly. A specific risk assessment process will be established which, as a minimum, will evaluate a proposal for development or adoption of an AI solution against the following criteria:

- Safety, Security and Robustness – no harm to people, things or environment
- Fairness and Bias – no unfair treatment based on who you are
- Accountability and Governance
- Transparency and Explainability – understanding of how AI works and decides.
- Contestability and Redress – ability to challenge AI decisions.
- Privacy and Information Security – CAA information is appropriately protected.
- Sustainability – the environmental impact of use of the technology is proportionate.
- Legal and Regulatory Compliance
- Use of Humans in the Loop – performance is monitored, and humans remain responsible for decision making.

4.3 This process will capture a risk assessment against each of these criteria, and where appropriate mitigations will be identified and put in place.

4.4 Residual risks will then be evaluated to determine whether a corporate risk should be raised under the standard CAA risk framework to be tracked and managed in the corporate risk register.

4.5 The use of this framework will be reviewed at least annually, and a determination made with regards to its continued use. Should the technology reach a point of stability and ubiquity such that a specialist risk assessment process is no longer considered necessary, the AI Risk Assessment Framework may be discontinued.

4.6 It should be noted that the use of the risk framework is intended for the development or adoption of bespoke or non-standard consumer AI applications, or the application of standard consumer AI applications to novel circumstances only. Use of the framework will not be required for the use of standard consumer AI applications in routine circumstances.

Chapter 5

Our Initial Flight Plan

5.1 We will identify and harness powerful AI use cases that propose significant enhancements to both colleague productivity and customer experience, while simultaneously allowing us to rigorously assess and refine our strategy for using AI.

5.2 We are committed to building a robust foundation for this AI-empowered future, developing essential tools like our AI Risk Framework, crafting adaptive policies and procedures, and delivering comprehensive training. Recognising that this transition represents a significant change for our colleagues and customers, we're committed to managing this evolution thoughtfully and effectively, ensuring everyone is supported through the journey.

5.3 This holistic approach ensures we are not just adopting AI but evolving with it – creating a CAA that's smarter, more responsive, and better equipped to lead in the dynamic world of aerospace.

The AI Portfolio

5.4 The activities and initiatives necessary to deliver against this strategy are far reaching across the organisation. This needs to be effectively managed and reported on according to our governance proposals.

5.5 To enable this, the AI Portfolio encompasses strategically significant activity from across the CAA that is associated with AI and autonomy. Managing this centrally enables the CAA to define specific outcomes, promote AI in aerospace and in the CAA, ensure a consistent approach that aligns to the strategy, efficiently monitor progress against the strategic plans, and provide transparent and coherent reporting to colleagues and customers.

The AI Strategy & Portfolio Hub

5.6 To ensure a consistent and comprehensive approach to regulating and utilising AI throughout the CAA, a dedicated AI Strategy & Portfolio Hub will be established as outlined in CAP3064, “The CAA’s Response to Emerging AI-Enabled Automation”.

5.7 Regarding the use of AI in the CAA, the AI Strategy & Portfolio Hub will play a crucial role in several areas:

- Providing leadership and strategic guidance on AI initiatives within the CAA.

- Promotion of AI throughout the organisation to foster a culture and mindset that proactively engages with AI.
- Overseeing and managing the AI Portfolio, encompassing the various programmes, projects, and initiatives related to the use of AI within the CAA.
- Ensuring that the CAA continues to value the application of the AI Principles in creating and maintaining trust in our use of AI.
- Collaborating closely with teams across the CAA to address the multidisciplinary challenges and opportunities when considering AI within our business and regulatory functions.

5.8 The AI Strategy & Portfolio Hub's oversight and expertise will also be instrumental in regulating AI in the aerospace sector. The Strategy for Regulating AI in Aerospace (CAP3064A) describes this in approach in detail.

Pre-Flight Checks (to Summer 2025)

5.9 Our flight plan for this journey starts with pre-flight checks – those actions that are necessary in the short term – followed by taxi and take-off where we can explore AI-enabled automation within the CAA from a robust and secure foundation of governance, understanding, and skills.

i) Establish an AI Risk Framework to enable colleagues and our leaders to understand the risks and benefits of AI early in their journey to explore, identify, and use AI tools safely throughout the business.

ii) Incorporate the strategy for using AI within our Customer, Consumer, People, and IT strategies, alongside key policy documents such as the CAA's IT Code of Conduct.

iii) Capture, assess, prioritise, and communicate internal use cases for AI, providing colleagues and leadership teams with examples of AI being applied to the CAA's work.

iv) Create a training development plan for all colleagues, responding to the risks identified and opportunities for upskilling colleagues.

v) Continue to provide colleagues with regular and accessible material and guidance, supplemented by internal communications activities.

vi) Establish the initial governance for AI, providing mechanisms to assure the CAA Board on our compliance against the 5 AI Principles.

Taxi & Take-off (to end of 2026)

5.10 Beyond our initial actions, we will be able to explore further use of AI, develop our colleague's understanding and capabilities further, and deliver practical uses of AI to enhance the CAA.

i) Develop a roadmap of pilot projects to provide the CAA with first-hand experience of applying AI to our workflows, to enable us to learn from early use cases, and to support the wider change management.

ii) Deliver several proofs of concept of AI applications in the CAA, giving affected colleagues and customers the opportunity to trial new solutions and enabling engagement across teams.

iii) Launch an AI Hackathon to support the roadmap for introducing AI-based tools across the business, supporting emerging AI skills in the CAA, and helping us to explore novel approaches to existing challenges.

iv) Using intelligence gained from proofs of concept and from research to continually inform our AI Portfolio and cross-pollinate with our Strategy for Regulating AI (Part A).

v) Develop the policies to govern AI architecture, adoption, development, use, and support across the CAA.

vi) Deliver policy updates for business functions including procurement, finance, human resources, to reflect the evolving use of AI, based on intelligence from the Taxi phase

Chapter 6

1. Getting involved, providing feedback and more information

6.1 We are committed to maintaining open channels of communication and collaboration as we implement and refine our AI strategy. We encourage engagement from all stakeholders and offer several ways to stay informed, provide input, and get involved:

CAA Website

6.2 Our website now includes a dedicated AI page at www.caa.co.uk/AI, serving as a central portal for all AI-related information. Here, you'll find the latest guidance, updates on regulatory developments, and links to relevant resources from across the CAA and beyond. We encourage you to visit this page regularly for the most up-to-date information on AI in aviation.

Email Contact

6.3 For enquiries, requests, or to share insights, please contact us at StrategyforAI@caa.co.uk. This central mailbox is monitored regularly, ensuring timely responses to your queries and contributions.

Feedback

6.4 We welcome feedback on all aspects of this strategy. Your input is invaluable in helping us refine our approach and ensure it remains relevant and effective. We aim to be collaborative, accessible, and supportive in enabling innovation in AI in aerospace and in the CAA, and your feedback plays a crucial role in achieving this goal.

A Note on AI Assistance

6.5 In the spirit of embracing the technologies we are discussing we would like to acknowledge that generative AI tools were used in the drafting of this strategy document. This approach has brought several benefits, including improved use of time, effective summarisation of complex information, and improved accessibility of the content. All content has been carefully reviewed and validated by our team to ensure accuracy and alignment with CAA policies and standards.

6.6 We look forward to working with you as we navigate the exciting and challenging landscape of AI in aviation. Together, we can ensure that the UK remains at the forefront of safe, innovative, and responsible AI adoption in the aerospace sector.

APPENDIX A

Abbreviations

AI	Artificial Intelligence
ATOL	Air Travel Organisers' Licensing
CAA	Civil Aviation Authority
CAP	Civil Aviation Publication
CIO	Chief Information Officer
GPT	Generative Pre-trained Transformer
HR	Human Resources
ICAO	International Civil Aviation Organisation
ISD	Information Services Department
IT	Information Technology
MOR	Mandatory Occurrence Report
NAA	National Aviation Authority
PDC	Performance & Development Conversation
SRO	Senior Responsible Owner
UK	United Kingdom

APPENDIX B

Illustrative Examples of AI-Enabled Automation within the CAA

To demonstrate the potential impact and versatility of artificial intelligence (AI) within CAA, we have compiled a range of illustrative examples. These examples span three key areas: regulatory functions, business operations, and everyday use. While not exhaustive, this list provides a glimpse into how AI could enhance our efficiency, decision-making processes, and overall effectiveness across various aspects of our work.

Regulatory Use of AI

1. Analysis of ATOL data to identify holders at higher risk of failure based on sales trends, pricing, and market changes.
2. AI-assisted initial evaluation of approval applications, summarising compliance and non-compliance areas to speed up processing, especially for large text submissions like operations manuals.
3. AI-powered analysis of digital Mandatory Occurrence Reports (MORs) to categorise, prioritise, and identify emerging themes, while also coding reports into appropriate systems.
4. Safety data analysis to detect subtle indicators ("weak signals") that may prompt investigation into potential safety or security risks.
5. Generation and analysis of regulatory risk scenarios, providing policy teams with insights to support policy development and rulemaking.

Business Use of AI

1. Support analysis of consultation responses by summarising overall reactions, identifying contentious points, and suggesting suitable response actions.
2. Evaluation of procurement bids to identify non-compliance and provide summary information for faster processing.
3. AI-powered chatbots to provide customers with guidance and information.
4. Creation of tailored content for different user groups based on common regulatory information.

5. Identifying necessary resource demands and possible changes based on demand trend data.
6. Accelerating processes for software development and testing within our information services department.

Everyday Use of AI

1. Summarising large documents to accelerate understanding and information processing.
2. Drafting initial versions of documents, images, videos, or audio based on notes or other sources to speed up work completion.
3. Streamlining the creation of Management Information reports.
4. Assessing and amending emails and document text to align better with cultures and language of international organisations or audiences.

APPENDIX C

Governance Proposals

This appendix proposes the governance structures, rules, and procedures that will be introduced, amended, and followed to support our responsible use of AI.

Management and ownership

The strategy and subsequent portfolio of activity for using AI within the CAA (Strategy Part B) is owned and overseen by the CAA's Chief Information Officer (CIO) as its Senior Responsible Owner (SRO). The policies governing the CAA's use of AI will be owned or overseen by the Information Serviced Department (ISD). Further detail relating to the responsibilities of individual elements are outlined below.

AI Portfolio

The proposal to establish an AI Portfolio to oversee and manage strategically significant AI activity across the CAA is described in CAP3064 "The CAA's Response to Emerging AI-Enabled Automation". Together with the Strategy and Portfolio Hub, this will incorporate relevant experts from across the business and support the cross-CAA functions and responsibilities outlined in this Strategy.

A clear business case for AI adoption

Any opportunity for the adoption of an AI solution must have a clear benefit for the CAA as a business, as a regulator, or for our customers, whether that be time or cost saving, customer experience enhancement, or improvements to the safety, security, or consumer value of aerospace. Any new technology carries an overhead to ensure it is supported and maintained effectively, and those overheads must be suitably offset by the benefits the solution brings.

Adoption of AI in the CAA

Enabling our colleagues with a human-first approach

All business processes utilising AI technology will be assessed regarding the appropriate degree of human intervention or oversight. We must be able to provide the assurance that decisions taken supported by AI are reviewed by humans where appropriate before meaningful action is taken.

Appropriate Solution Design

Any solution must conform with the CAA technical reference model, best practices for software development, and application vulnerability management. Solution designs must

appropriately manage security, privacy, safety, legal, fairness, and sustainability risks through the governance and risk frameworks described in this strategy.

Considered service support and change management

We will develop and refine suitable practices for the ongoing support and development of new AI systems. Production and deployment of AI systems will follow standard CAA change management practices, with the required approvals at each stage, and effective communication tailored as necessary for users and other stakeholders.

AI System Design Principles

All bespoke AI solutions will be developed in line with a set of specific design principles. These principles will be established and maintained by the Architecture and Design team in ISD, with support and oversight from the AI Centre of Excellence, and as a minimum ensure any intended use is clearly understood and documented, suitable operational metrics and performance monitoring information is captured, data quality requirements are clearly understood and mechanisms are in place to maintain that data quality, and there is an appropriate level of transparency and explainability in place. Any new solution will be approved by the Architecture and Design team as compliant with the design principles.

AI System Ownership

All bespoke AI solutions and services will have a designated business owner responsible for management of risk and utilisation of the solution. They will be responsible for ensuring the solution is managed appropriately, and decommissioning the solution when it is no longer fit for purpose or of continuing business value.

AI System Owners will carry out periodic reviews of the solution to ensure they remain in line with the risk assessment established for the system. The owner of the system will be responsible for responding to questions and challenges that may arise relating to the system they own.

With delegation from the SRO, the overall responsibility for use of AI at the CAA will sit with the Head of the AI Centre of Excellence, ensuring that solutions are developed, enhanced, and reported in line with the CAA's AI Framework. The Centre of Excellence, AI Community of Practice, and ISD will be equipped to provide AI System Owners with guidance and expertise. These groups will also coordinate the development and implementation of training, templates, and procedures for AI System Owners.

AI Risk Assessment

A risk assessment in line the CAA AI Risk Framework at Chapter 4 will be carried out for any new AI solution or when functionality is significantly modified, to ensure the solution remains within the CAA risk appetite. This assessment may be carried out as part of the business case analysis (new solutions) or by the system owner (existing solutions with significant modification). The risk assessment must be reviewed at least annually by the AI

system owner to ensure it remains current and accurate in the light of changing features of the solution or context of its use.

AI System Repository

We will establish and maintain a repository of the bespoke AI systems and services in use at any given time. The repository will capture the purpose, ownership, business criticality, and review history for the system/service. This applies only to bespoke solutions developed for the CAA, not to consumer AI tools in more general use. This does not exclude the capture of suitable information regarding these systems in both EasyVista and iServer, our configuration management database and enterprise architecture tools respectively. Any system in use by the CAA is captured in these repositories and that will also be true of any AI systems.

AI Citizen Development

The existing process to identify, evaluate and monitor solutions developed by non-IT colleagues which are business critical will also apply to those developed using AI tools. Where the solution goes beyond basic use of a consumer-focused AI solution and represents a unique business solution developed using such tools, it will be logged as a citizen developer solution, and will be subject to the same risk management and oversight processes as other such solutions.

AI introduced to existing software

It is anticipated that a range of service providers will start to introduce AI features into existing software in use by the CAA. The most obvious example of this is the Microsoft Office tooling for which Copilot has been introduced. At present this is a separate paid for add-on, but certain AI features are already making their way into the main office products, and it is likely that more will follow in future.

Introduction of any new features into software in use by the CAA can carry a degree of risk. Uncontrolled rollout of new features may represent a security risk for example, either through vulnerabilities in the new features themselves or vulnerabilities caused by inappropriate use of such features.

The CAA currently evaluates new features being added to existing software to determine if any such features need to be configured or blocked to mitigate such risks. The same approach will be taken with regards to AI features. ISD are responsible for this evaluation and any mitigation actions that are required.