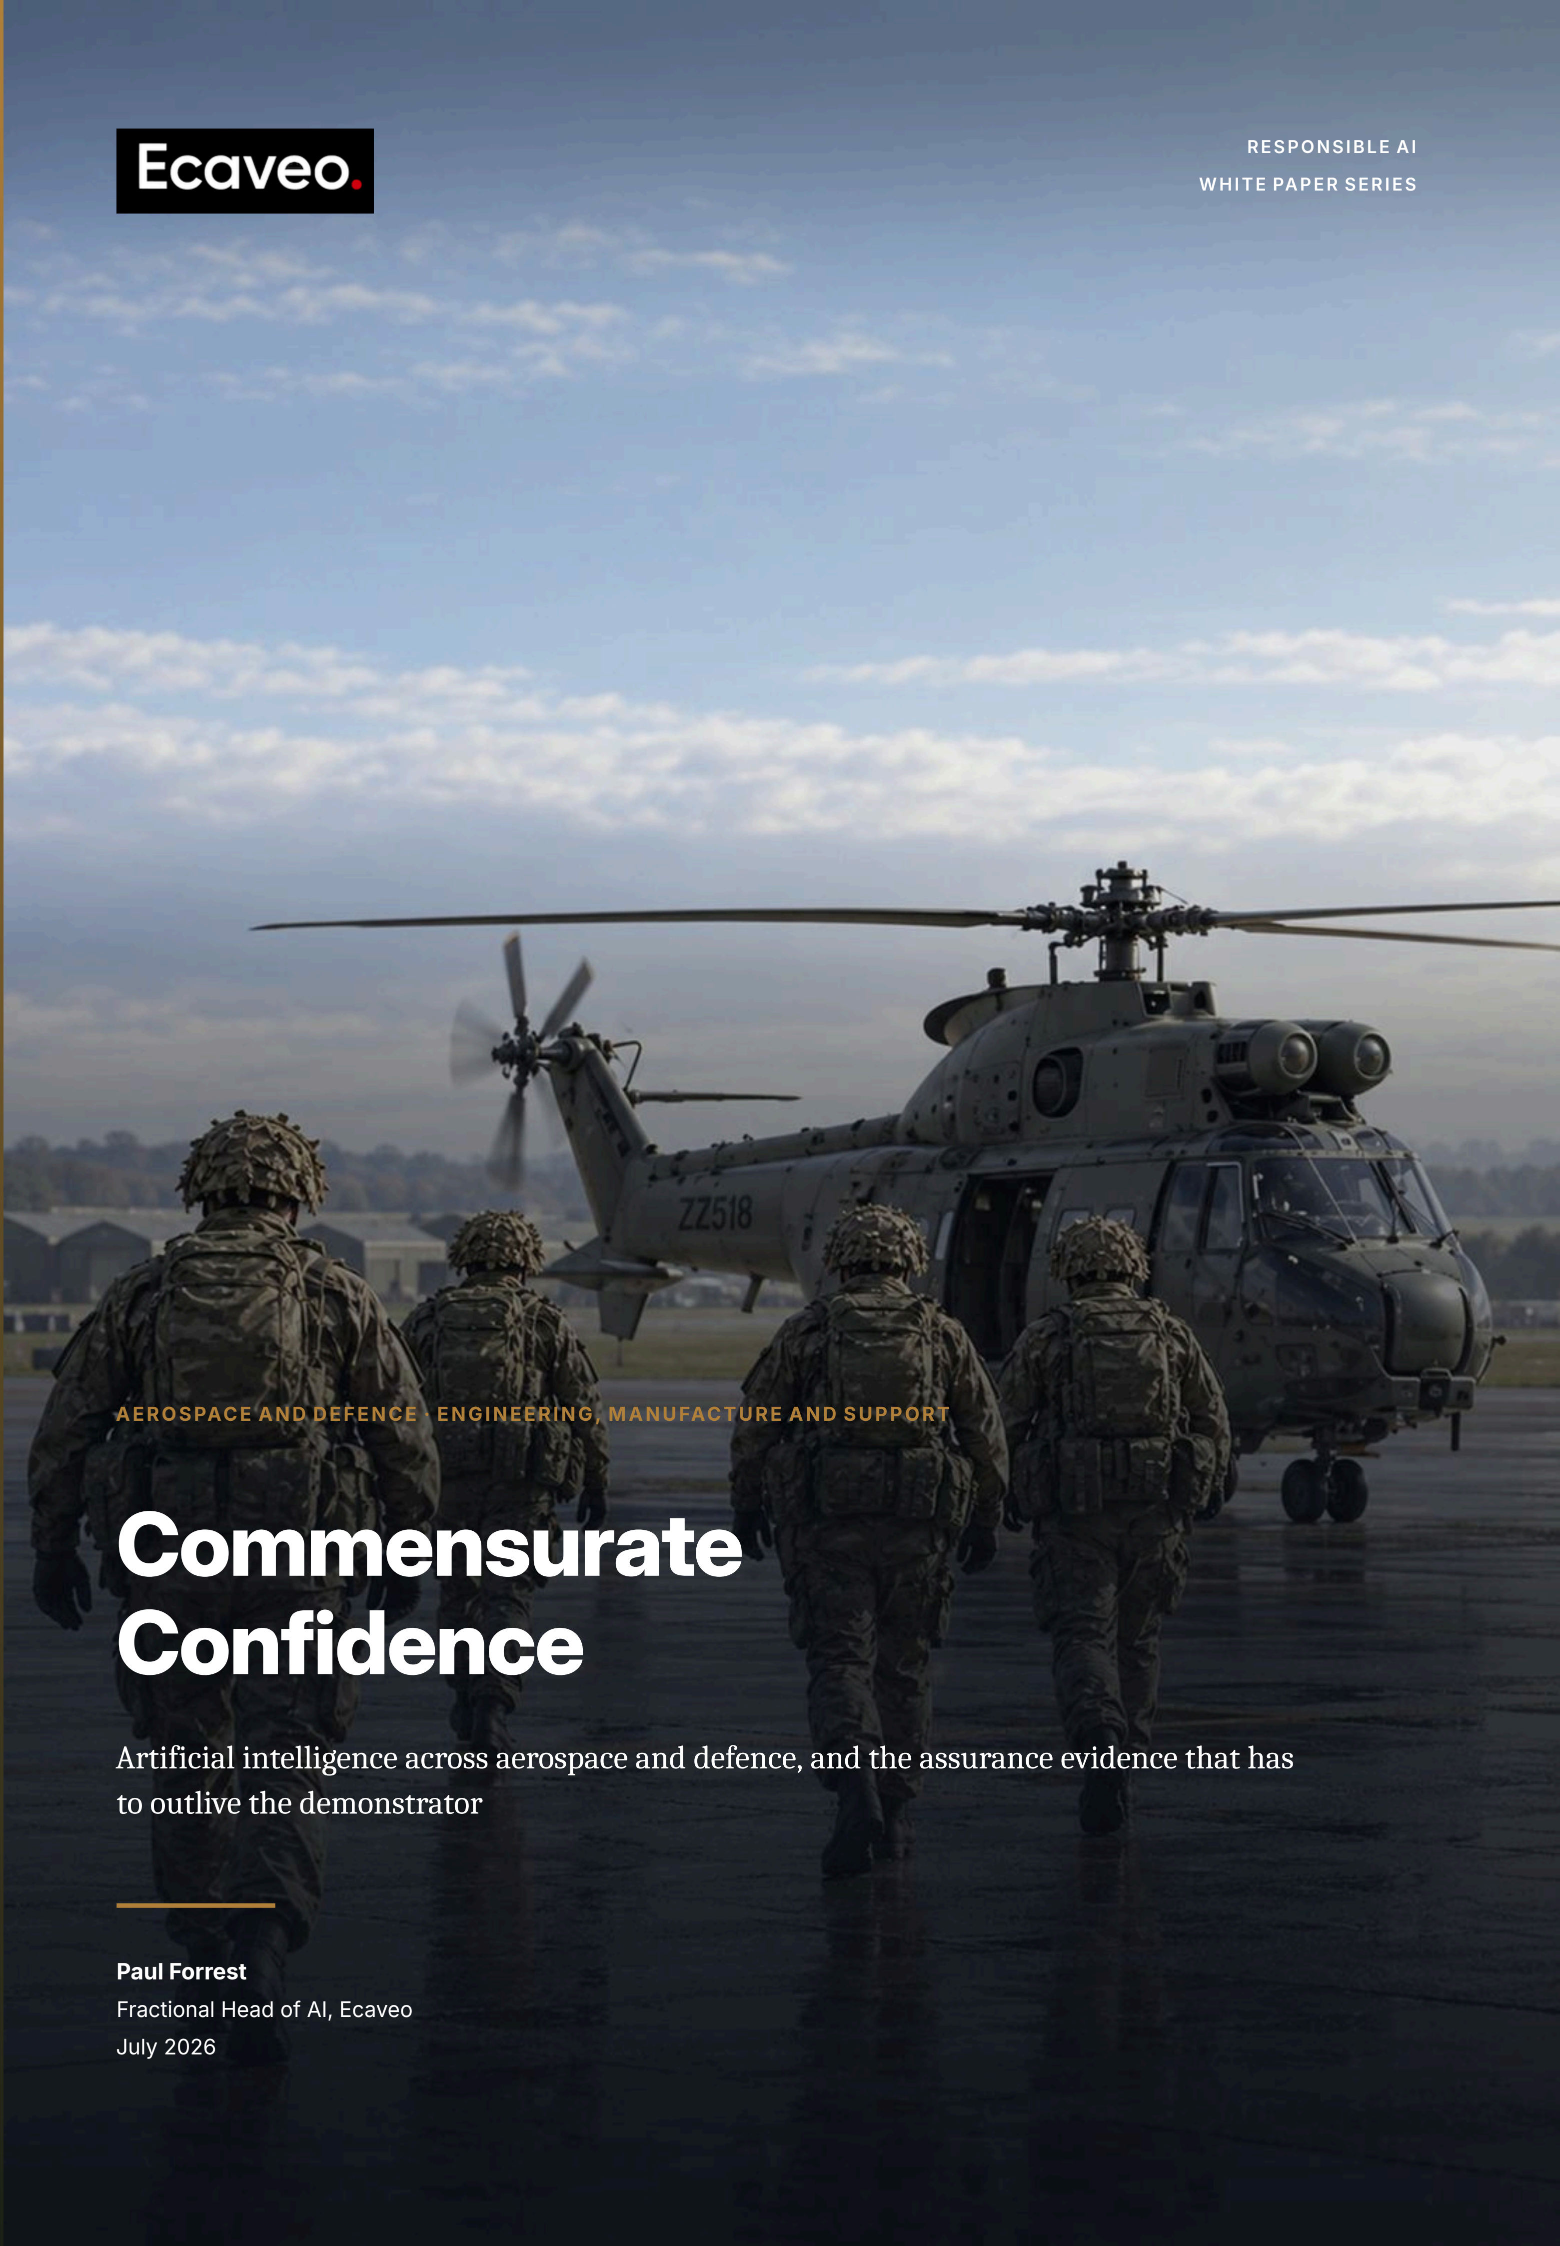

Ecaveo.

RESPONSIBLE AI
WHITE PAPER SERIES

AEROSPACE AND DEFENCE · ENGINEERING, MANUFACTURE AND SUPPORT

Commensurate Confidence

Artificial intelligence across aerospace and defence, and the assurance evidence that has to outlive the demonstrator

Paul Forrest
Fractional Head of AI, Ecaveo
July 2026

About this paper

This paper covers aerospace and defence companies, meaning industry, though the customer through most of it is Defence. It sits alongside papers on financial services, public services and critical national infrastructure, freight operations, and a trilogy on professional firms. Each stands alone.

The title comes from JSP 936, which declines to offer a dependability scale and states instead that the level of confidence developed in the system must be commensurate with the reliance placed upon it.[5] STANDARD That sentence does more work than any framework diagram. It sets a variable standard, tied to consequence, and it puts the burden of establishing where the standard sits on the people building the thing.

Editorial balance. Roughly seven parts in ten of this document is evidence and sector analysis, two parts in ten is original framework, and one part in ten concerns how I would run the work. On the value of predictive maintenance in particular, the honest evidence position is weaker than the sales material and Chapter 5 sets out why.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence supports it.

REGULATOR A regulator, government department, audit body or official statistics publisher speaking in its own voice.

STANDARD A published standard, statute, defence standard or formal policy document.

CASE LAW An accident investigation finding, judgment or enforcement action.

PEER-REVIEWED A primary study, systematic review or parliamentary investigation working from primary material.

SURVEY A survey or dataset with disclosed method and sample.

VENDOR RESEARCH Research from a body with a commercial or membership interest in the answer. Directional only.

PRACTICE Practitioner opinion, including my own.

EVIDENCE GAP A question the available evidence does not answer.

ECAVEO My own framework, proposed rather than proven.

Legal, regulatory and safety notice. This paper describes defence policy, airworthiness regulation, certification standards, export control and safety obligations throughout, and Chapters 2, 3 and 6 in particular. Nothing in this document constitutes legal advice, and nothing in it substitutes for a safety case, an airworthiness determination or an export control classification. The provision of legal advice sits outside the terms of any engagement with the author or with Ecaveo, and the material is presented to support discussion and further review by qualified advisers.

Statements about what a source says are attributed to that source. Where I draw a practical implication the source does not itself state, the text says so. Several standards described here were in draft when this was written.

Contents

THE ARGUMENT IN ONE PARAGRAPH

JSP 936 binds the Ministry of Defence and does not bind its suppliers. The obligation reaches industry only where a contracting team writes it into a contract, which means a company waiting for the clause will be assembling evidence retrospectively about work it has already done. In civil aerospace the position is starker still, because no published certifiable process standard for machine learning exists, the joint EUROCAE and SAE standard remains in draft with a target of end 2026, and its first issue will be capped at design assurance level C and will exclude adaptive learning. This paper argues that the assurance case is the product. A company that builds the evidence chain as it engineers, from need through configuration and test into service, can sell into both markets from one controlled record. A company that builds demonstrators will keep building them.

ILLUSTRATIVE MATERIAL

Every company, programme, platform and person in the worked examples is invented, and each is marked as an illustrative scenario where it appears. Published research, official statistics, accident findings and standards are cited and numbered.

CITATION

Forrest, P. (2026). *Commensurate Confidence: Artificial Intelligence in Aerospace and Defence.* Ecaveo Responsible AI Series, Aerospace and Defence. Version 1.0, July 2026.

ALSO IN THIS SERIES

The Known Safe State. Public services.

Partial Understanding. Financial services.

The Protected Constraint. Operations.

Verified Authority. Law firms.

A NOTE FROM THE AUTHOR

On evidence that has to last 30 years

Every other sector in this series worries about whether its evidence will survive a regulator. This one has to build evidence that survives the aircraft.

My background is business process work from the 1990s, then machine learning and language models from 2014, and I build small specific models in preference to large general ones. That preference has never been easier to defend than in this sector, where a system with a narrow, stated operating domain is assurable and a general one is a research problem wearing a product's clothes.

What makes aerospace and defence different from everything else I have written about is the timescale. A commercial deployment that turns out badly gets switched off. A platform stays in service for decades, through configuration changes, obsolescence management, capability insertions and several generations of the people who built it. Evidence produced today has to be intelligible to somebody in 2050 who has never met anybody involved.

That is not a new problem for this industry, which has been managing configuration and safety evidence at that timescale for a very long time and is better at it than any other sector I work in. The new part is that machine learning breaks some of the assumptions the existing apparatus rests on. Requirements traceability assumes requirements can be written down. Structural coverage assumes structure. Regression testing assumes you know what the previous behaviour was.

There is one more thing I want to flag at the start, because it shapes how the rest of this reads. JSP 936 contains a sentence in its cover note that I have not seen matched anywhere in government. It says that in many ways the document represents the aiming point or ideal end state, as many of the supporting tools and frameworks are necessarily still in development, and that where the requirements cannot yet be fully met, the risk should be understood and a route to compliance planned.[5] STANDARD That is a remarkable admission for a directive to make about itself, and it is more useful to a supplier than any amount of confident guidance would be.

TWO PAGES THAT CIRCULATE ON THEIR OWN

Executive summary

The customer's policy does not bind you, the certification standard does not exist yet, and the most instrumented fleet in history is getting less available. Everything else follows from those three.

1

JSP 936 binds the Ministry of Defence and not its suppliers. The directive must be applied by all MOD staff across the system lifecycle. Its general statement about industry is that meeting the requirements should be a shared endeavour between MOD and its suppliers.[5] STANDARD Obligations reach a company through contract, where MOD teams contracting for AI must require suppliers to demonstrate competence to deliver products to a level of confidence commensurate with their use.

2

There is no published process standard for machine learning in civil aviation. ED-324, also designated ARP6983, is the joint EUROCAE and SAE process standard for developing and certifying safety-related aeronautical products implementing AI. It remains in draft, with EUROCAE listing a target publication date of 31 December 2026. Issue 1 is scoped to non-adaptive supervised machine learning up to design assurance level C, and excludes information security and human factors.[12] STANDARD

3

The regulator has published a ceiling. The EASA concept paper issued for comment in June 2026 caps supervised learning in initial airworthiness at design assurance level C, and states within its treatment of net safety benefit that allocation of level D or the equivalent lower levels to an AI constituent involved in catastrophic failure conditions will not be accepted.[10] STANDARD Both are published and available to plan against now.

4

Availability on the most data-rich fleet ever built has fallen. The United States Government Accountability Office reported in June 2026 that the F-35 mission capable rate fell from 67 per cent in financial year 2021 to 44 per cent in financial year 2025, with full mission capability down from 38 per cent to 25 per cent.[22] REGULATOR That is the counterweight to any claim that instrumentation and analytics produce availability on their own.

5

The economic case for predictive maintenance rests on simulation. The published cost-benefit work on prognostics and condition-based maintenance in aviation is discrete-event simulation and life-cycle cost modelling.[23,24] PEER-REVIEWED I located no measured, independently verified improvement in UK military aircraft or vehicle availability attributable to AI-based predictive maintenance. EVIDENCE GAP

6

The demand signal from the customer is weak, and the Defence Committee said so. A Commons Defence Committee report of 10 January 2025 identified a say-do gap, noting that speeches by senior leadership frequently refer to AI but are rarely accompanied by specifics about what MOD wants to achieve, how, and by when.[21] PEER-REVIEWED It valued the UK defence AI sector at approximately £285m in 2023, projected to £1.2bn by 2028.

What follows for a chief engineer

Build the assurance case as you engineer, and not afterwards. The evidence JSP 936 expects covers data provenance from collection to ingestion, configuration management of all output-influencing data, verification and validation across the operating domain including boundary conditions and realistic edge cases, and demonstrated safe behaviour outside that domain.[5] STANDARD None of that can be reconstructed from a working demonstrator.

Define the operating domain before the model. An operational design domain that is written down, bounded and testable is the single artefact that both the defence directive and the civil concept paper depend on, and it is the one most often left implicit because writing it down is uncomfortable.

Assume the standard will land at level C and plan the architecture accordingly. Issue 1 of ED-324 introduces the machine learning constituent as the transition point from conventional system development into a distinct development process, deliberately leaving the existing DO-178C and ARP4754B ecosystem intact.[12] STANDARD Architecting to that boundary now costs little and converts later.

Chapter 9 sets out an 18-month sequence from demonstrator to fielded capability. Chapter 10 states what would change the position taken here.

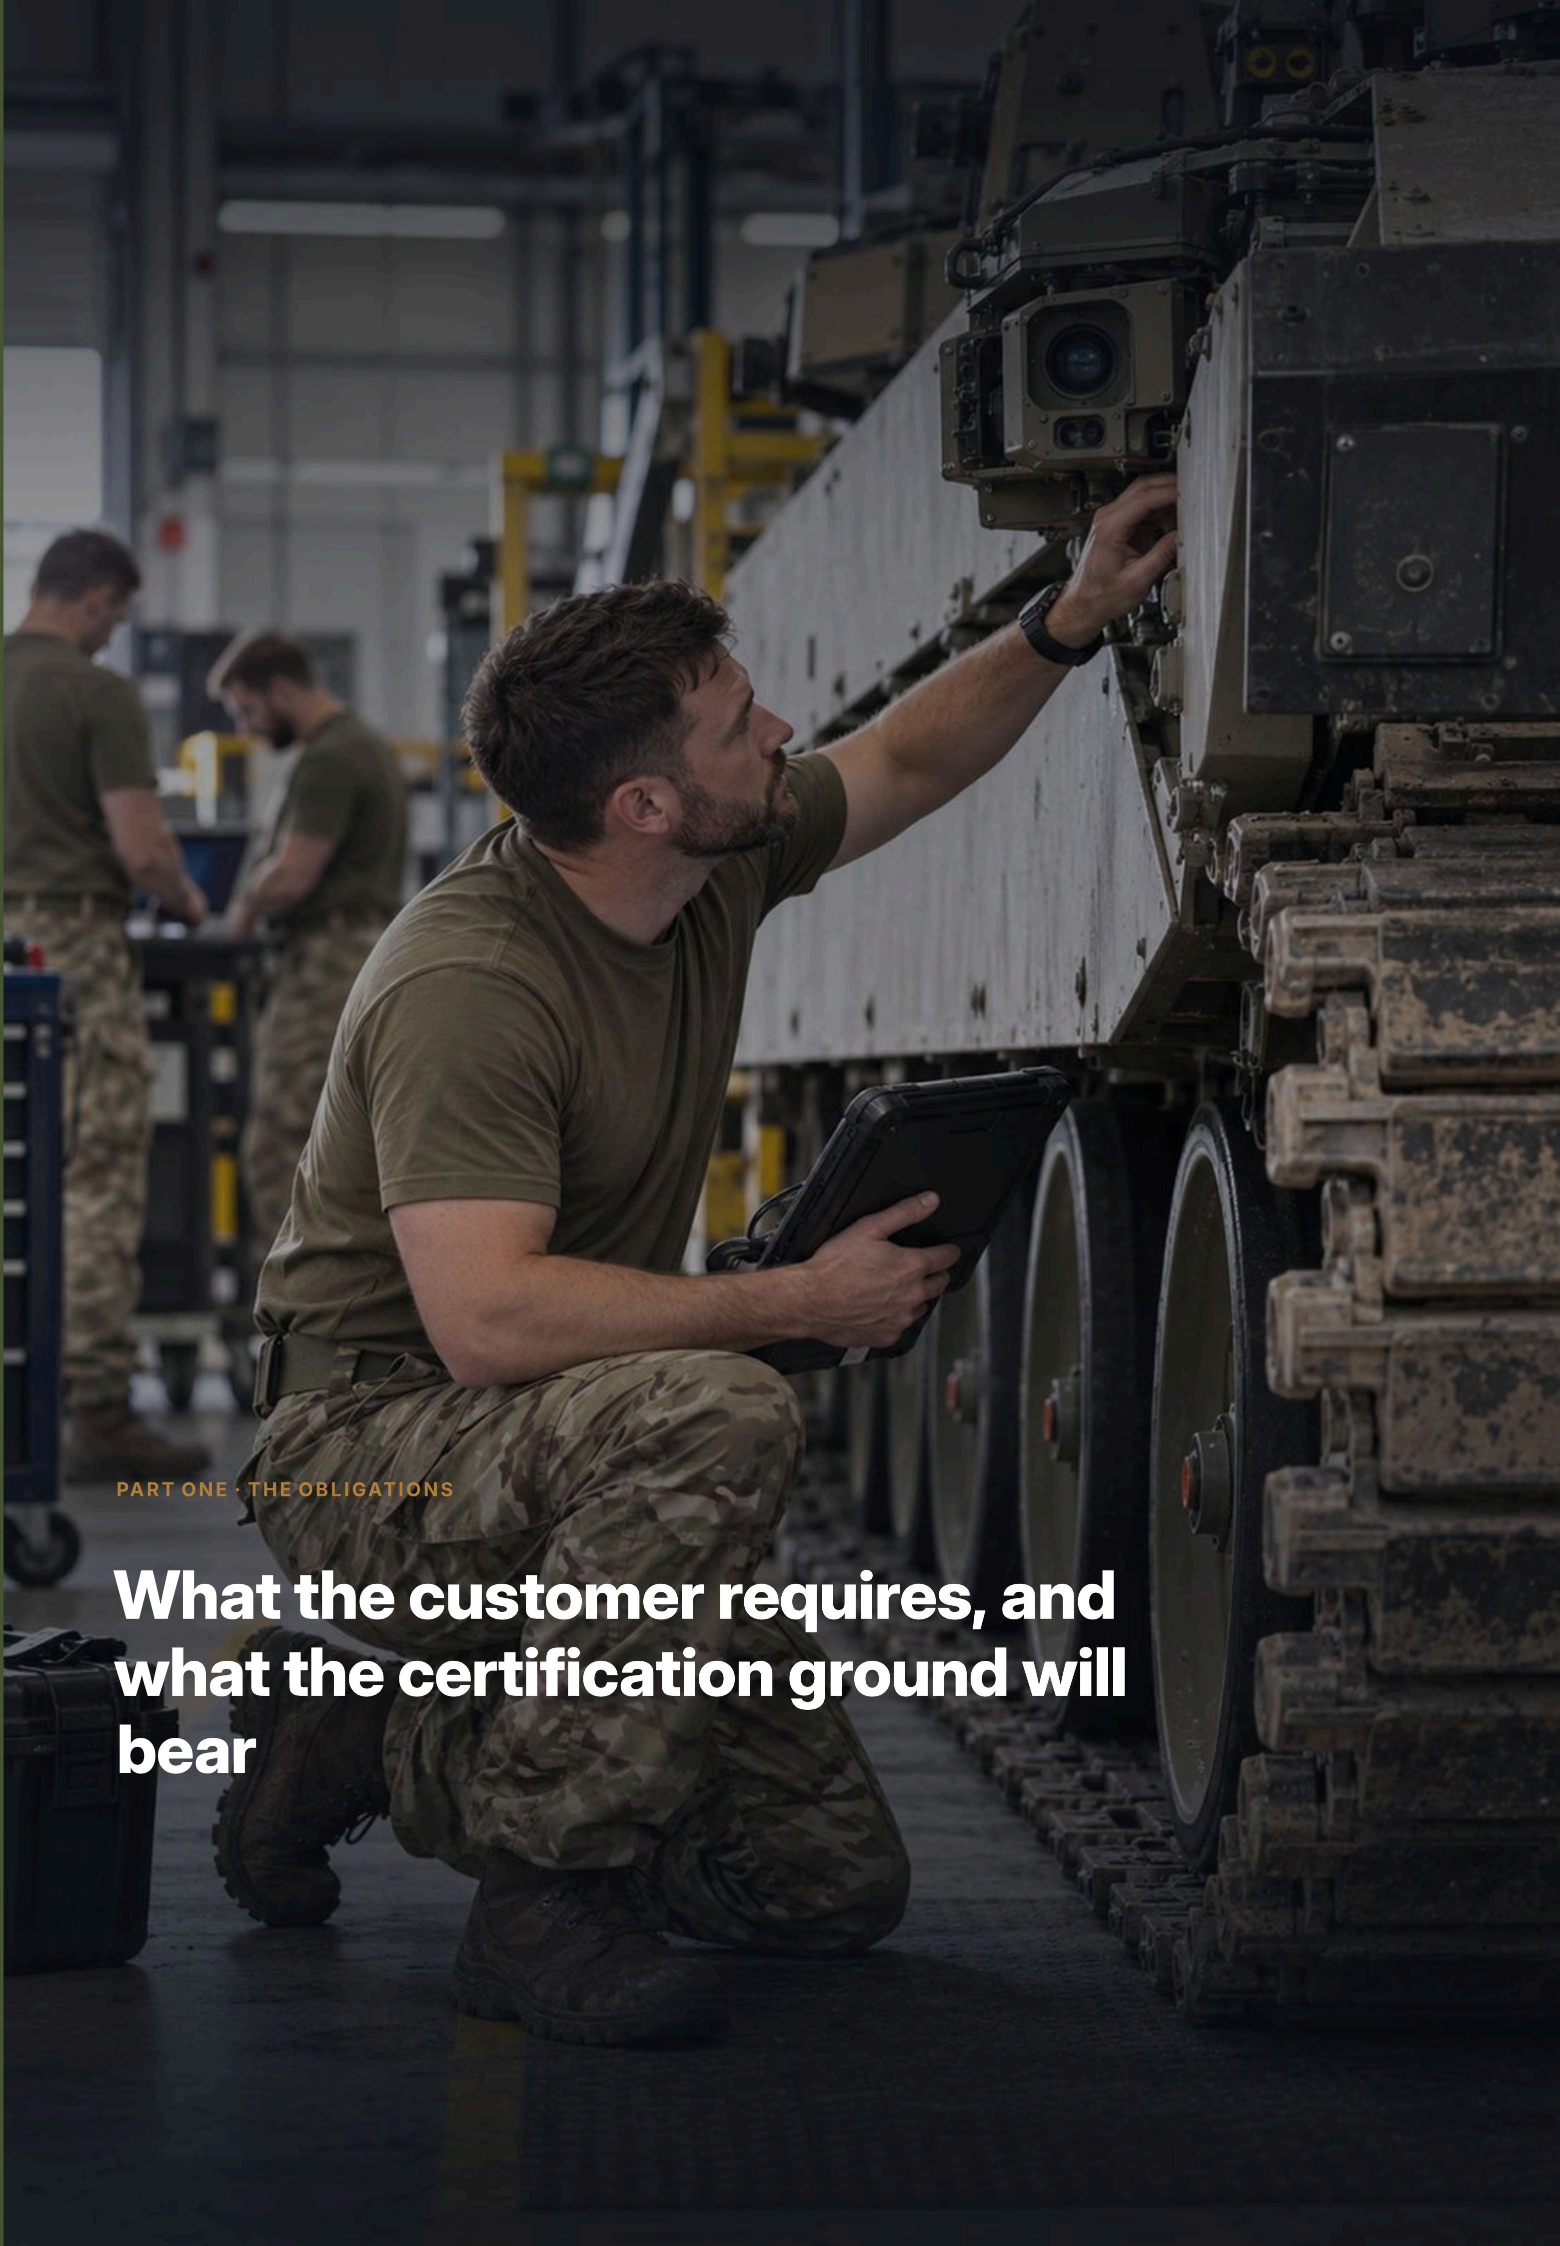

PART ONE · THE OBLIGATIONS

What the customer requires, and what the certification ground will bear

CHAPTER ONE

Dependability is the strategy

In most sectors assurance is a cost applied to a capability. Here it is the thing being sold, and companies that understand that price their work differently.

The UK aerospace, defence, security and space sectors together generated an estimated £46.8bn of gross value added in 2025, on turnover of nearly £110bn, employing 468,500 people directly and exporting more than £50bn, which is 45 per cent of sector turnover.[20] VENDOR RESEARCH Aerospace alone accounts for £41.5bn of turnover, £14.3bn of value added, £30.6bn of exports and 113,000 employees, at productivity of £126,400 per worker. Defence reported turnover of £36.5bn, with exports of £10.8bn stated as a five-year average. Whole-sector productivity is given as £99,900 of output per worker, described as nearly a quarter above the UK average. These are trade association figures compiled from member data, which makes them the best available and not independent.

The customer side is larger and moving. The Ministry of Defence budget stood at £62.0bn for 2025/26 with £68.3bn planned for 2026/27, and NATO-qualifying spend at 2.3 per cent of gross domestic product in 2025 rising to an estimated 2.6 per cent in 2026, against commitments of 3.5 per cent of GDP on core defence by 2035 plus 1.5 per cent on wider defence and security-related investment.[19] REGULATOR Of £40.6bn paid by the MOD core department to industry in 2024/25, 45 per cent went out non-competitively, and 49 per cent of new contracts by number were awarded on a single-source basis, described as the highest level of single sourcing since 2015/16.[18] REGULATOR

I mention the single-sourcing figure because it bears on how assurance obligations actually travel. In a competitive tender a requirement can be specified and priced. In a single-source negotiation it is argued about, and the party with the better evidence position wins the argument.

1.1 Why lifecycle evidence is the differentiator

Consider what the customer's own directive expects of the material a supplier hands over. Data provenance must be assured and recorded from the point of collection to ingestion by the AI. All data influencing output, including metadata, must be kept under configuration management controls commensurate with those applied to the AI itself. Verification and validation must demonstrate behaviour across the operational design domain including boundary conditions and realistic edge cases, and safe behaviours outside the domain must be demonstrated. Where the AI is updated, testing must cover both existing and new behaviours, through new test cases and regression testing, to give assurance that intended outcomes were achieved and that unwanted emergent effects such as catastrophic forgetting and model drift have not occurred.[5] STANDARD

Read that as a specification for a records system and the commercial position becomes clear. A company that has generated those records as a by-product of ordinary engineering can respond to a requirement in weeks. A company that has a working demonstrator and no records has to either reproduce the development under configuration control or argue that its evidence is adequate, and both cost more than doing it properly the first time.

THE COMMERCIAL POINT

Every one of those requirements is cheap during development and expensive to reconstruct. A supplier who has them is holding an asset a competitor cannot buy.

The directive also states something that a supplier should quote back in the right circumstances. Where the differences between AI and conventional software present too much assurance risk, it says, in some applications these issues will simply present too much risk and in such cases AI must not be adopted.[5] STANDARD That is a licence to decline, issued by the customer, and it is worth having in the file when a programme office is pressing for a capability the evidence will not support.

1.2 The substitution argument, and its limits

Buried in the assurance section of JSP 936 is the argument that most programmes will eventually rely on. Where AI is replacing extant technology or human decision-making, the assurance case may be able to argue that the specific risk from the AI is no greater than the system being replaced.[5] STANDARD

That is a sound principle and it is harder to discharge than it looks. Making the comparison requires a characterised baseline for the incumbent, meaning the error rate of the human or the legacy system, measured under the same conditions, on the same distribution of inputs. Most organisations do not have that. Nobody measured how often the experienced maintainer got it wrong, because there was never a reason to.

My own practice is to treat baseline characterisation as the first deliverable of any programme intending to use the substitution argument, and to budget for it separately. PRACTICE It is unglamorous, it takes a quarter, and without it the argument is an assertion.

CHAPTER TWO

What the customer actually requires

The single most common overclaim in this sector is that JSP 936 applies to suppliers. It does not, and understanding exactly how the obligation travels is worth real money.

JSP 936, Dependable Artificial Intelligence in Defence, Part 1 Directive, was published on 13 November 2024 at version 1.1 and has not been substantively revised.[5] STANDARD Part 2, promised at paragraph 100 to cover the roles and responsibilities of the Responsible AI Senior Officer and the delegated responsibilities in the governance chain, has not been published.

Its scope is broad. Software is defined to include program code, operating procedures, all relevant data and documentation. Examples given as in scope include object detection, recognition and identification in intelligence and surveillance, reinforcement learning in command and control and in autonomous systems, large language models in human resources applications, and AI in decision support toolsets covering course of action development, task allocation across crewed and uncrewed force mixes, mission tactics engines and logistics planning. Out of scope are openly available search engines, predictive text on commercial messaging, and common office productivity tools where the risk of the resulting material is low, with responsible use of those remaining the user's responsibility.[5] STANDARD

2.1 Who it binds, precisely

The requirements must be implemented from the outset of all Defence Capability development where MOD requirements lead to the use of AI, and the directive must be applied by all MOD staff in all phases of the system lifecycle, from pre-concept through to disposal or service termination.[5] STANDARD The list of those who must apply it runs through science and technology providers at all technology readiness levels, customers and senior responsible owners, requirements managers, delivery agents, defence lines of development owners, trials units, specialist engineering functions, policy makers and end users.

Every one of those is inside the Ministry. The directive's general statement about industry is at paragraph 14, and it is hortatory. Paragraph 199 adds that suppliers new to Defence must understand the operational context, information and operational security, and Ministry policy, which is a statement about industry and not an obligation on it. Meeting the JSP requirements should be a shared endeavour between MOD and its suppliers.[5] STANDARD

The obligation therefore travels through Section 8, which addresses contracting teams. Externally acquired AI must attract the same level of confidence that the requirements have been met as AI developed within or for the MOD, and the directive acknowledges that this can be more challenging than for traditional software, with MOD teams possibly having to stand up additional assurance capabilities to address evidence shortfalls. Teams contracting suppliers of AI must require them to demonstrate competence to deliver products to a level of confidence that is commensurate with their use, and where a machine learning model is to be used in a safety-related application the supplier must be required to demonstrate appropriate people, processes, experience and evidence to support the safety argument. Where third-party or open-source software is used, suppliers must be required to develop compelling arguments as to why confidence in the resulting AI has not been undermined.[5] STANDARD

Two further contracting provisions matter commercially and are easy to miss. Contracting conditions should consider through-life support including access to data and algorithms as well as the resulting models, with clear ownership of intellectual property for data, algorithms, models and tests. And assessment of restrictions caused by foreign export controls should include data, as well as algorithms and models, to ensure that MOD has the necessary freedom of action to maintain, modify, upgrade and operate the AI.[5] STANDARD

> **Two terms to get right, because a specialist reader will check.** JSP 936 does not use the term "responsibility gap". It uses "accountability gap", once, at paragraph 125(b), requiring analysis of the allocation of functions between human and AI agents to ensure that there is no accountability gap, meaning that humans remain accountable for and in control of the effects of the system. Nor does the directive use "TEV&V". It uses verification and validation, at paragraphs 169 to 173.[5] STANDARD Both terms are common in United States and research usage and neither should be attributed to this document.

2.2 The governance chain a supplier will meet

Each top-level budget must appoint a Responsible AI Senior Officer, who must ensure that appropriate assurance for the ethical and responsible development and use of AI is in place across all AI-based projects under their accountability. For AI not developed in the UK, the officer must satisfy themselves that the AI and the data meet UK policy and provide a Statement of AI Ethics Assurance. An AI ethical risk assessment addressing the ethical principles and potential harms must be used at project outset and on material change.[5] STANDARD

The five ethical principles come from the 2022 policy paper on delivery of AI-enabled capability, and their names are worth knowing precisely, since they appear in assessments a supplier will be asked to support. They are Human-Centricity, Responsibility, Understanding, Bias and Harm Mitigation, and Reliability.[7] STANDARD

How mature is any of this? The Ministry published a report on 3 October 2025 recording 17 nominated Responsible AI Senior Officers across Defence commands and component organisations, and describing implementation as at a formative stage with momentum building.[8] REGULATOR It identifies gaps in suitably qualified and experienced personnel, a need for dedicated support teams, insufficient senior leadership training, and difficulty translating high-level policy into use-case-specific guidance. On industry it observes that some suppliers are taking an exceptionally conscientious approach, alongside persistent challenges of infrastructure limitations, unclear development pathways and inconsistent supplier capability, and states an intention to set clearer expectations through contract terms and procurement aligned to secure by design.

There are no quantitative programme data in that report. It is a qualitative maturity assessment, and it should be read as one.

2.3 The demand signal, and the committee that named the problem

The Commons Defence Committee reported on developing AI capacity and expertise in UK defence on 10 January 2025. It valued the UK defence AI sector at approximately £285m in 2023, projected to reach £1.2bn by 2028, and noted that 75 to 80 per cent of defence AI companies are small enterprises or start-ups. It observed that the United States and China each spend more than four times what the UK government spends on AI.[21] PEER-REVIEWED

Its central criticism is the one suppliers repeat privately. The Committee found a say-do gap, noting that although speeches by senior leadership frequently refer to AI, these are rarely accompanied by specifics about what MOD wants to achieve, how, and by when. It found the demand signal weak, concluding that a clear pipeline was needed to attract capital investment and help industry understand the skills it would need. It found strong experimentation and poor scale-up to fielded capability, and it identified barriers for non-traditional suppliers including procurement complexity, security clearance delays, classified data access and slow decisions.[21] PEER-REVIEWED

Some of that has since been answered. A letter to Defence personnel from the Defence Secretary and the senior leadership team, published on 10 June 2026, announced a Rapid AI Delivery Taskforce focused on four areas covering machine-augmented intelligence fusion, operational advantage in degraded environments through real-time electromagnetic understanding, AI-driven planning automation, and AI-enabled autonomous swarms for contested environments.[6] REGULATOR It also commits to a federated Defence data infrastructure across all classification levels and to modular procurement avoiding vendor lock-in. Note what it is, though. It is correspondence, not a strategy or a policy instrument, and it contains no funding figures. Any pound figure attached to it in a business case is an inference.

CHAPTER THREE

Where civil certification stands

The European regulator has published a great deal and mandated almost nothing. The standard everyone is waiting for is at draft seven with 1,800 comments against it.

The European Union Aviation Safety Agency published its Artificial Intelligence Roadmap 2.0 in May 2023, and it remains current. Nothing labelled version 3.0 exists, and material that appears to be a later roadmap is generally the concept paper or the rulemaking proposal, both of which are deliverables under Roadmap 2.0.[9] STANDARD

The roadmap's four trustworthiness building blocks are worth learning because the vocabulary recurs across the sector. AI trustworthiness analysis characterises the application and covers safety, security and ethics-based assessment, and the roadmap describes it as a gate to the other three blocks. AI assurance covers learning assurance, which addresses the shift from programming to learning because existing development assurance methods are not adapted to learning processes, together with development and post-operations explainability. Human factors for AI covers operational explainability and human-AI teaming. AI safety risk mitigation addresses the residual risk arising where the box cannot be opened to the extent required.[9] STANDARD

3.1 The classification everyone will be asked about

The roadmap's classification of AI applications is the framework a customer will use to place a proposed capability, so a supplier should be able to locate its product in it without hesitation.

EASA CLASSIFICATION OF AI APPLICATIONS

LEVEL	SUB-LEVEL	DESCRIPTION
Level 1, assistance to human	1A	Human augmentation
Level 1, assistance to human	1B	Human cognitive assistance in decision-making and action selection
Level 2, human-AI teaming	2A	Human and AI-based system cooperation
Level 2, human-AI teaming	2B	Human and AI-based system collaboration
Level 3, advanced automation	3A	The AI-based system performs decisions and actions that are overridable by the human
Level 3, advanced automation	3B	The AI-based system performs non-overridable decisions and actions, for example to support safety on loss of human oversight

TABLE 1 From Roadmap 2.0, figure 4. The concept paper issued for comment in June 2026 extends coverage down to a level 0 for low automation and up to level 3.

The roadmap's own timeline expects finalised guidance for level 1 and 2 applications in 2026, finalisation of overall AI and machine learning policy in 2028, and adaptation to further innovation in 2029. Its stakeholder prognostic, which it is careful to describe as a prediction and not a formal target, puts first approvals of level 1 applications at 2025, first approvals of level 2 and 3A at 2035, and autonomous AI at 2050 and beyond.[9] STANDARD Anybody building a business case on certified airborne autonomy inside a decade is arguing against the regulator's own published expectation.

3.2 What has actually been published, and what has not

Concept paper issue 02, guidance for level 1 and 2 machine learning applications, was published on 6 March 2024. A proposed issue 03, guidance for safety-related artificial intelligence applications, was released for a 10-week comment period on 3 June 2026 and extends coverage across levels 0 to 3.[10] STANDARD It carries the criticality ceiling quoted in the executive summary, which is that for supervised learning in initial airworthiness the ceiling is design assurance level C. Separately, in its treatment of the net safety benefit concept, it states that allocation of level D or the equivalent lower levels to an AI constituent involved in catastrophic failure conditions will not be accepted, and that reduction below that is not acceptable either. The second statement constrains the one-level reduction a net safety benefit credit can buy, and it is narrower than the first.

The Agency's first regulatory proposal on AI, NPA 2025-07, opened for a three-month consultation on 10 November 2025, aligning AI trustworthiness requirements with the high-risk regime of the EU AI Act.[11] STANDARD That is a proposal. It is not a rule.

> **No published certifiable process standard for machine learning exists in civil aviation.** ED-324, designated ARP6983 by SAE, is being produced jointly by EUROCAE working group 114 and SAE G-34. EUROCAE's own deliverables table lists it as draft with a target publication date of 31 December 2026. A joint status presentation of August 2025 described it as at draft seven with approximately 1,800 first-ballot comments being dispositioned. Issue 1 is scoped to non-adaptive machine learning in supervised mode, up to design assurance level C, excluding information security and human factors, and applies to airborne and air traffic management, crewed and uncrewed.[12] STANDARD Any claim that an AI-specific aerospace certification standard exists today is wrong.

The architectural idea inside the draft is the part worth acting on now. It introduces the machine learning constituent as the transition point from conventional system development and safety processes into a distinct development process, which deliberately leaves the existing ecosystem intact around it.[12] STANDARD A company that isolates its learned components behind a defined constituent boundary today will map onto the standard when it lands. One that has distributed learned behaviour through a conventional architecture will not.

3.3 The standards the AI has to live inside

None of the conventional apparatus has gone away and two important documents have been superseded, which catches people out.

CERTIFICATION STANDARDS, WITH STATUS AS AT JULY 2026

STANDARD	TITLE	STATUS
DO-178C and ED-12C	Software considerations in airborne systems and equipment certification	Published 2011, current, no AI supplement
ARP4754B	Guidelines for development of civil aircraft and systems	Published December 2023, superseding ARP4754A of 2010
ARP4761A	Guidelines and methods for conducting the safety assessment process on civil airborne systems and equipment	Published December 2023, superseding ARP4761 of 1996
ED-324 and ARP6983	Process standard for development and certification of aeronautical safety-related products implementing AI	Draft, EUROCAE target 31 December 2026

TABLE 2 The two December 2023 supersessions are the ones most often missed. A document citing ARP4754A as current is nearly three years out of date.

3.4 The two regulators a UK company answers to, and the gap between them

The UK Civil Aviation Authority does have an AI strategy. CAP3064 and its parts describe the Authority's response to emerging AI-enabled automation, with CAP3064A covering the strategy for regulating AI and advanced automation in aerospace, at version 1.3.2 dated 3 December 2024. CAP3019, an AI technology outlook, is dated 12 August 2024.[13] STANDARD The Authority describes itself as both an AI-enabled and an AI-enabling regulator, and states an objective of enabling the UK aviation sector to adopt AI with confidence.

Both remain at version 1.3.2, unchanged for more than 18 months, during which the European Agency moved from a concept paper to a rulemaking proposal.[11,13] STANDARD The series itself has grown, with a Part C and a Part G on consumer and AI principles added since, so the Authority has not been idle. What it has not done is revise the regulatory strategy. My reading, which the Authority would presumably contest, is that a UK company serving both markets should treat the European material as the de facto specification and the UK material as a statement of regulatory posture. PRACTICE

On the military side there is a cleaner gap. The Military Aviation Authority's regulatory publications cover general, flying, air traffic management, continuing airworthiness engineering and type airworthiness engineering, with more than 15 manuals alongside. Regulatory Article 1600 at issue 10 categorises uncrewed air systems as open, specific or certified, listing capability for autonomous operations as an optional attribute in all three categories, and it defines autonomous operations and artificial intelligence for those purposes.[14] STANDARD I could find no dedicated regulatory article, manual or position paper on AI or machine learning assurance. EVIDENCE GAP UK military aviation therefore has no published AI-specific airworthiness regulation, and JSP 936 is not an airworthiness instrument. That absence is itself worth stating to a programme board, because it means the airworthiness argument will be constructed case by case against Def Stan 00-055 issue 5 and Def Stan 00-056 issue 8.[15] STANDARD

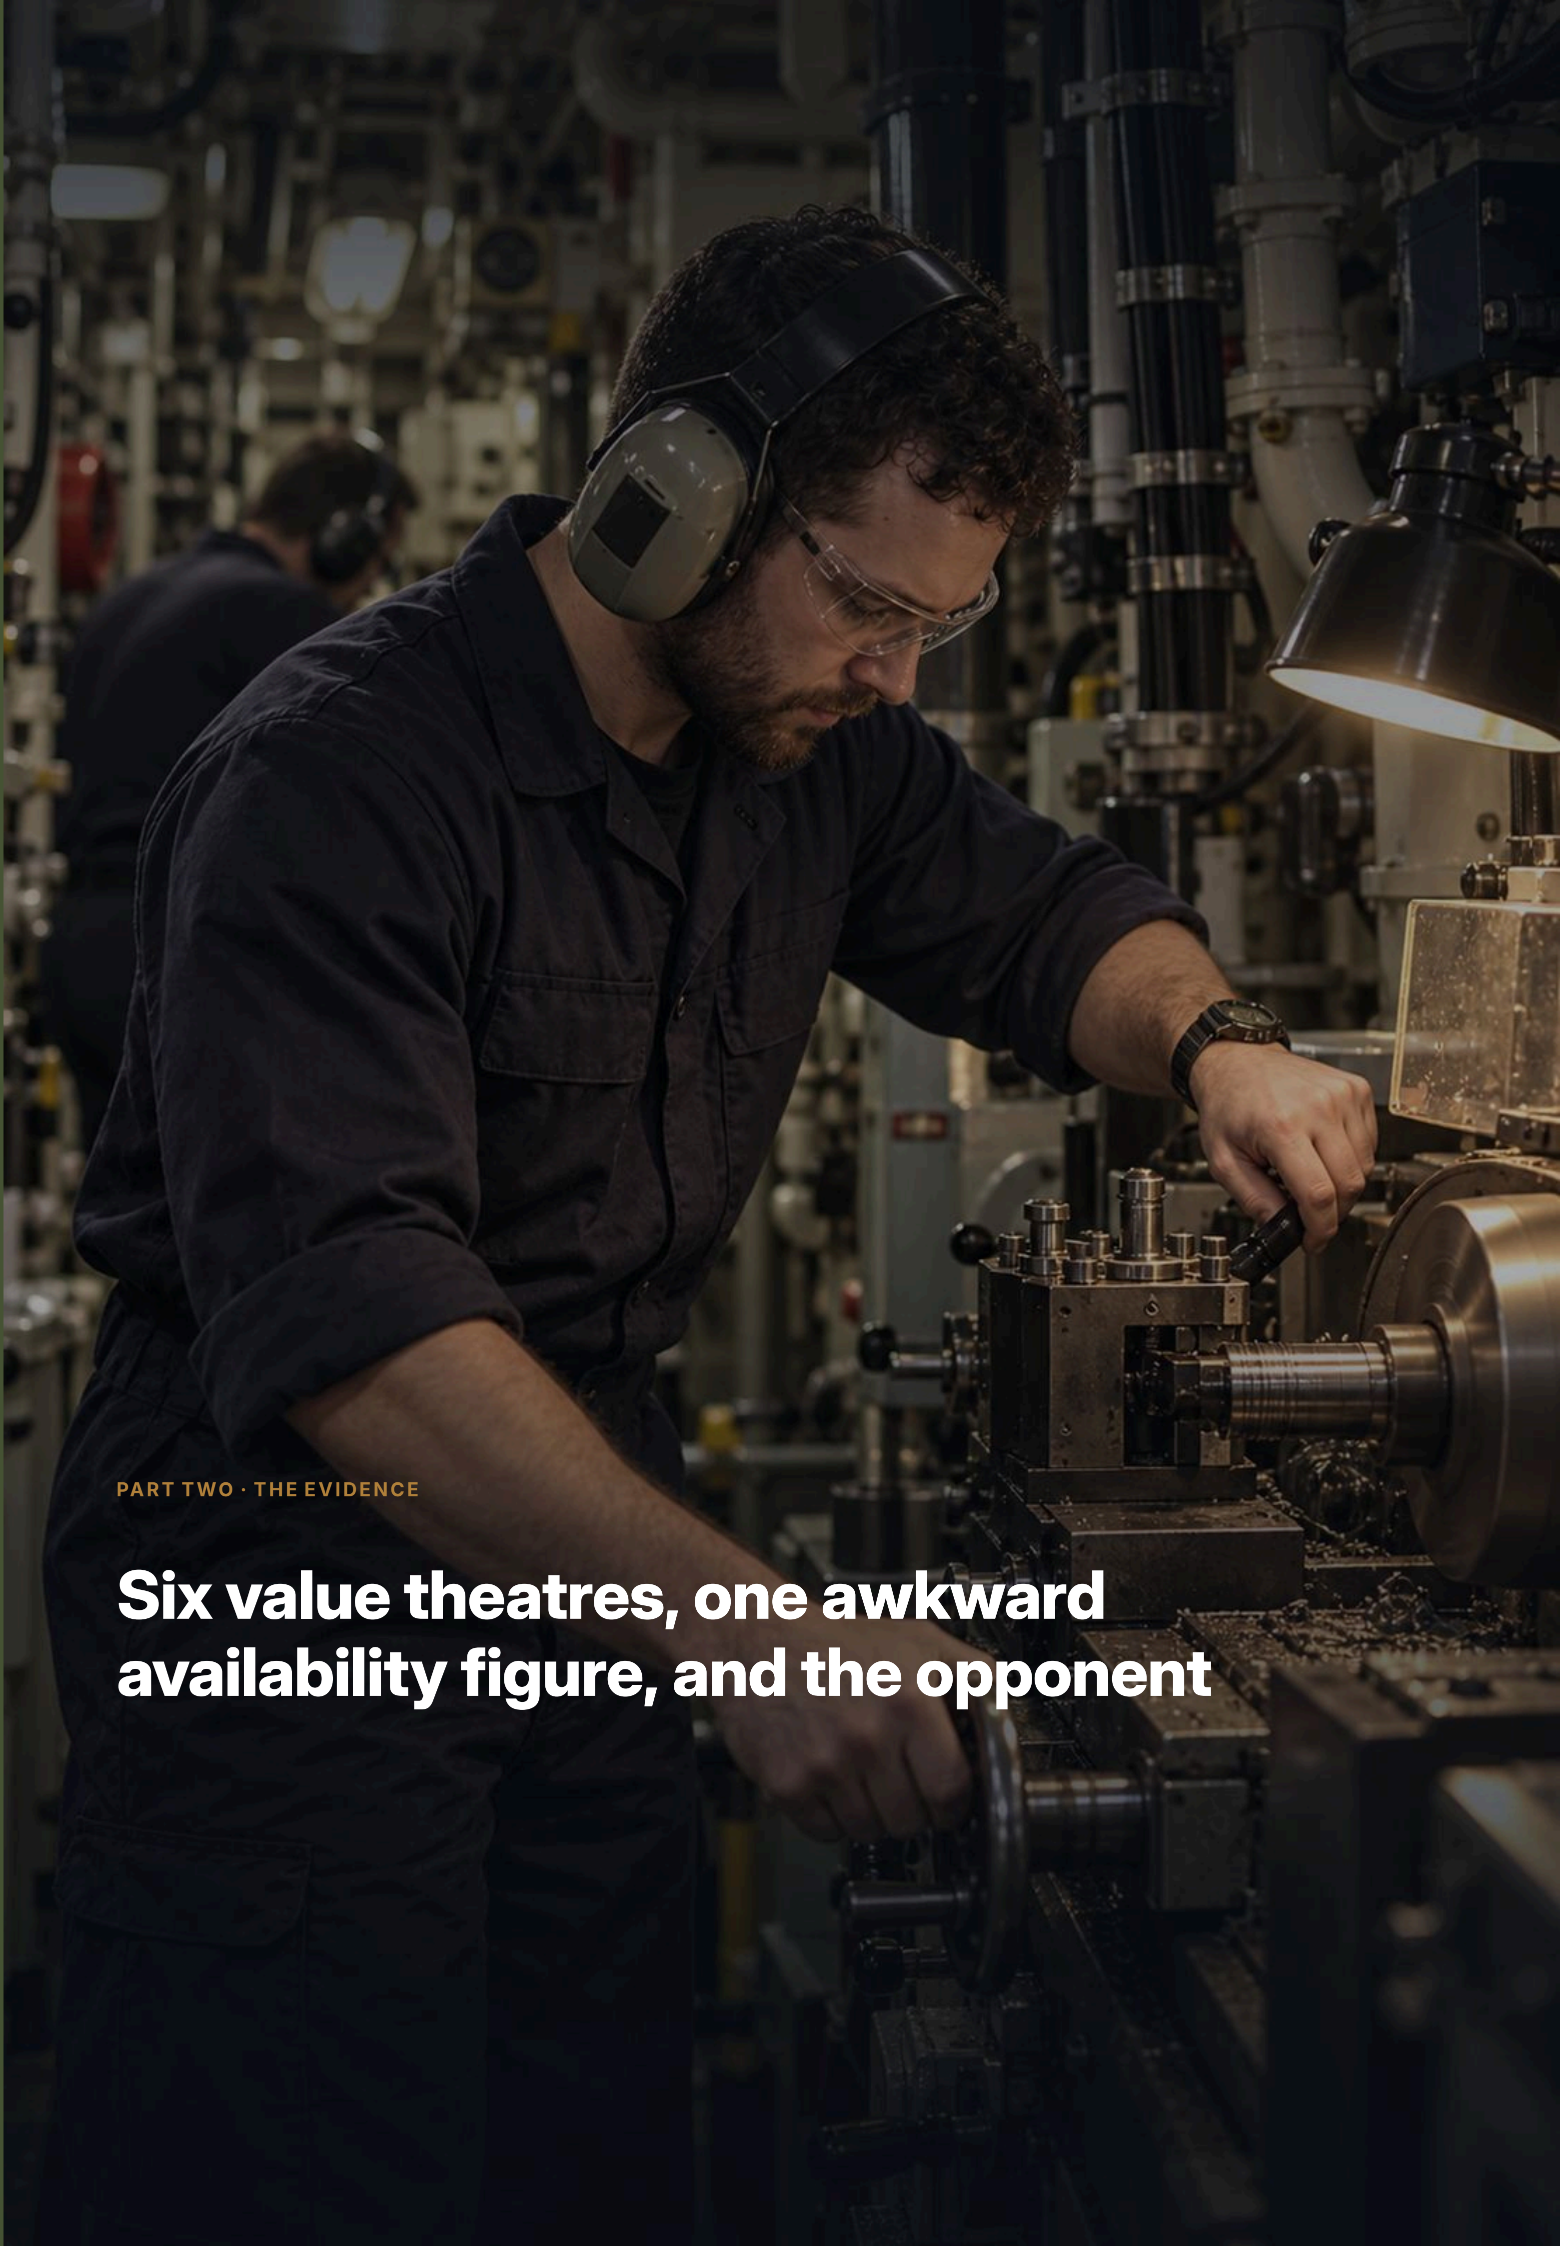

PART TWO · THE EVIDENCE

Six value theatres, one awkward availability figure, and the opponent

CHAPTER FOUR

Where the value is credible

Six theatres, ordered by how confidently the value can be shown. Enterprise workflows come first, which is not where the interesting engineering is.

THE SECTOR PORTFOLIO

THEATRE	OPPORTUNITY	VALUE TEST
Enterprise productivity	Bids, programme controls, knowledge retrieval and secure software work.	Cycle time, quality against review, information leakage rate and total delivery cost.
Engineering and the digital thread	Retrieve requirements, compare configurations, detect inconsistencies and support design review.	Review time, requirement coverage, defect escape rate and traceability at audit.
Supply chain and obsolescence	Map dependencies, forecast constraints and evaluate substitution evidence.	Lead time, shortage exposure, approved alternatives found and supplier resilience.
Manufacturing quality	Inspect components, analyse process variation and direct skilled review.	Detection by defect type, false reject rate, scrap, rework and line disruption.
Maintenance and availability	Predict degradation, prepare work packs and retrieve approved technical information.	Availability, unscheduled removals, maintenance error and evidence completeness. See Chapter 5.
Mission and operational support	Fuse bounded information, surface uncertainty and support planning under human command.	Decision time, missed signals, calibration and performance when degraded.

TABLE 3 Six value theatres. The value test states what would have to improve against a measured baseline. Where the published evidence is weak, the entry says so.

4.1 The evidence on engineering work, which points both ways

Two randomised trials of AI assistance on high-skilled engineering work reach opposite conclusions, and a serious paper has to cite both.

The positive result comes from three randomised controlled trials embedded in ordinary business operations at Microsoft, Accenture and a third large firm, with 4,867 developers randomised to access an AI coding assistant, finding a 26.08 per cent increase in completed tasks with a standard error of 10.3 per cent, and larger gains among less experienced developers.[25] PEER-REVIEWED The conflicts should be disclosed. The assistant studied is a Microsoft product, and two of the three sites are the vendor and a large systems integrator.

The negative result comes from a randomised trial with task-level allocation, in which 16 experienced open-source developers worked 246 tasks on mature repositories where they averaged five years of prior experience. Allowing AI increased completion time by 19 per cent.[26] PEER-REVIEWED The sample is very small, the domain is a single one with unusually high quality bars, and the tooling is from early 2025, so the result should not be generalised. The finding inside it that does travel is the misperception. Developers forecast a 24 per cent speedup, and afterwards still estimated they had been 20 per cent faster.

Reconcile the two and the picture is consistent with everything else in this series. Assistance helps most where the person is least expert and the task is most routine, and it helps least, or hurts, where the person is highly expert and the task sits in territory the model has not learned well. An aerospace engineering organisation has a great deal of the second kind of work.

4.2 Requirements, where the industrial evidence is missing

A systematic literature review of large language models for requirements engineering covered 74 primary studies from 2023 and 2024, finding a shift from defect detection towards elicitation and validation, predominantly using GPT-based models with zero-shot and few-shot prompting.[27] PEER-REVIEWED Its stated limitation is the one that matters here. Research predominantly occurs in controlled environments with limited use in industry settings, and there is an industrial validation gap with limited evaluation inside real development workflows.

There is domain-specific work. A paper presented at the Foundations of Software Engineering conference in June 2025 evaluates large language models for requirements question answering in industrial aerospace software.[28] PEER-REVIEWED I have cited it because it exists and is directly on point, and I should say that I could not obtain its full text, so I am not relying on its numbers. EVIDENCE GAP

The practical reading is that requirements traceability review is a promising application with a real evaluation problem, since the ground truth is a known-defect set that most organisations do not maintain. Building that set is the first task, and it is worth doing whether or not the model ever arrives.

4.3 Manufacturing inspection, and the benchmark trap

Deep learning for surface defect detection is a mature research field with a substantial published literature, including systematic reviews in the major applied AI journals.[29] PEER-REVIEWED Headline accuracy figures in that literature are almost always reported on curated public benchmarks.

Those benchmarks are not a prediction of yield on an aerospace production line. The line has rare-defect class imbalance, a changing part mix, lighting that varies with the shift, and an inspector population whose disposition decisions are themselves the ground truth being learned from. Cross-dataset benchmarking studies, which test a model trained on one dataset against another, are where generalisation failure becomes visible, and they are much less commonly cited in vendor material.[30] PEER-REVIEWED That inference about the gap between benchmark and line is mine. PRACTICE

Shadow mode on one inspection point answers the question directly and cheaply, which is why it appears as a pilot in Chapter 8.

CHAPTER FIVE

The availability question

The strongest argument for AI in this sector is that it will improve availability. The strongest available evidence points the other way, and the honest response is to say so and then explain why.

5.1 The F-35 figure

The United States Government Accountability Office reported in June 2026 on F-35 sustainment. The mission capable rate fell from 67 per cent in financial year 2021 to 44 per cent in financial year 2025. The full mission capable rate fell from 38 per cent to 25 per cent over the same period. The Global Support Solution reset is estimated to need $13.7bn more than planned through financial year 2031, lifetime United States sustainment cost was estimated at $1.6 trillion in 2024, and more than $7bn of additional parts are needed from the private sector.[22] REGULATOR The Office also records that the Department of Defense paid its contractor hundreds of millions in incentives since 2020 to improve F-35 readiness, and that these incentives have not been effective.

Two cautions before anybody uses this. It is a United States report about a United States programme, and UK figures are not published in comparable form. And nothing in it attributes the decline to a failure of analytics, because the causes it identifies are parts availability, depot capacity, the reset programme and contractual incentives.

The inference is mine and I will state it plainly. The F-35 is the most heavily instrumented combat aircraft ever fielded, generating health and usage data at a volume no previous programme approached. Over four years, on that fleet, mission capability fell by 23 percentage points. Whatever the causal story, this is not a fleet where data richness has translated into availability, and any business case asserting that it will elsewhere needs to explain what is different. PRACTICE

WHY THIS MATTERS HERE

No aircraft in history generates more health and usage data. If instrumentation and analytics produced availability by themselves, this fleet would show it.

5.2 What the published economics of predictive maintenance actually rest on

The academic case for prognostics and condition-based maintenance in aviation is real, careful and almost entirely modelled.

A widely cited cost-benefit analysis for commercial aircraft uses discrete-event simulation of a 150-seat short and medium range aircraft under a condition-based maintenance programme, and it does something most such work does not, which is model prognostic errors explicitly.[23] PEER-REVIEWED It is a simulation study. An earlier life-cycle cost study of prognostic health management for helicopter avionics is likewise a modelling exercise.[24] PEER-REVIEWED A systematic literature review of predictive maintenance for aircraft concluded that no comprehensive survey of applications and techniques devoted to the aircraft manufacturing industry existed, that research focus is biased towards engines because of the lack of publicly available datasets, and that greater automation is needed to optimise aircraft maintenance.[31] PEER-REVIEWED

The customer's own showcase carries no numbers either. The Defence AI Playbook's spare parts failure prediction case study, covering British Army in-service vehicles some dating back to the 1980s, states its benefits qualitatively and flags its own data problem, noting that legacy vehicles lack modern digital health and usage monitoring systems, so the work may depend on manually recorded data and service records with consequent quality issues around accuracy and consistency.[16] REGULATOR

So the defensible statement is this. I located no measured, independently verified improvement in UK military aircraft or vehicle availability attributable to AI-based predictive maintenance. The published economic case rests on simulation and life-cycle cost modelling. The one large contemporaneous data point on a data-rich fleet shows mission capability falling. EVIDENCE GAP

5.3 Why I still think the theatre is worth entering

Nothing above says predictive maintenance does not work. It says nobody has demonstrated that it does, at fleet scale, in a way that survives scrutiny. Those are different claims and the difference is the whole of the opportunity.

A company that runs a properly designed evaluation, with a characterised baseline, a held-out failure set and an independent analyst, will hold something genuinely scarce, which is evidence. Given the state of the literature, a single well-executed field evaluation would be publishable, and it would be worth more in a single-source negotiation than a decade of case studies.

The obstacle is the structural one. Safety-critical fleets replace components before failure, which produces a severe shortage of run-to-failure data, and that shortage is permanent.[32] PEER-REVIEWED Approaches that work around it, including physics-informed models, accelerated life testing and synthetic augmentation with validated failure physics, are where I would put engineering effort, and I would be sceptical of any supplier whose method assumes the failure data exists.

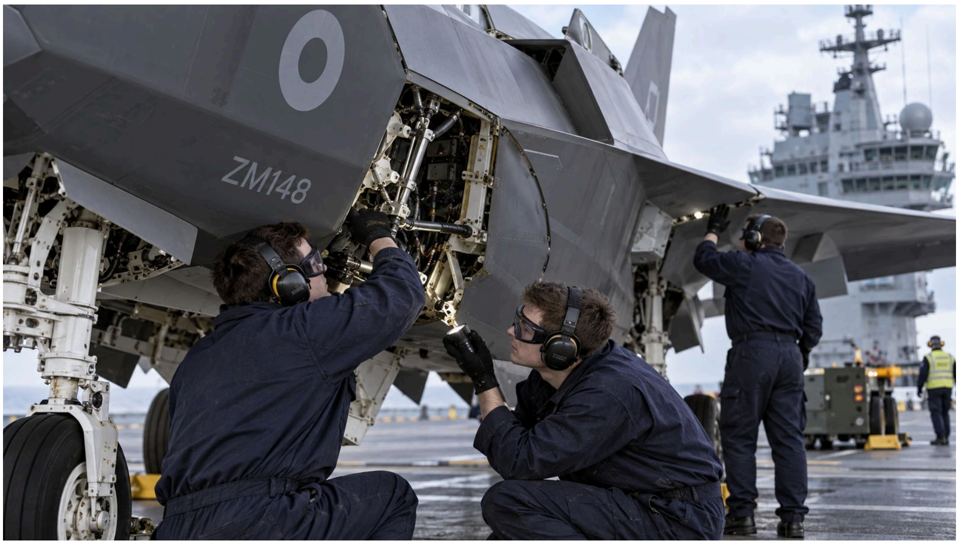

Carrier engineering. Maintenance evidence, configuration state and the technical publication set are the three things a maintainer needs to match, and matching them is a retrieval problem before it is a prediction problem.

CHAPTER SIX

Adversarial conditions and controlled information

Every other sector in this series designs against error. This one designs against an opponent who is trying to cause the error.

6.1 What adversarial machine learning has actually been shown to do

The canonical demonstration remains the physical-world attack on visual classification, where black and white stickers applied to a stop sign produced targeted misclassification in 100 per cent of laboratory images and 84.8 per cent of captured video frames from a moving vehicle.[33] PEER-REVIEWED It is a 2018 result on a road sign, and it establishes the principle that a physically realisable, low-cost perturbation can defeat a trained classifier.

Closer to this sector, a review of adversarial robustness in aerial detection covers optical, infrared and synthetic aperture radar sensing across digital and physical attacks. Results it surveys include physical optical attacks reducing target detection scores across a 25 to 85 per cent span with success rates between 60 and 100 per cent, attack success rates of 75 and 78 per cent in another study, adversarial cloud attacks reducing F-measure from 0.8253 to 0.2572, infrared attack success exceeding 90 per cent using thermal-insulation and temperature-control materials, a separate infrared implementation costing under $5, and black-box attacks on synthetic aperture radar succeeding 98.5 per cent of the time.[34] PEER-REVIEWED

The review's own limitations deserve equal weight with its numbers. It records insufficient real-world testing, no standardised defence evaluation framework, and reliance on low-resolution datasets rather than operational-grade imagery.[34] PEER-REVIEWED

> **An honest statement about what is not documented.** I found no verified, officially reported case of an adversarial machine learning attack against a fielded defence system. Claims circulating about drone spoofing, camouflage patches against military sensors and similar were not traceable to an official finding. EVIDENCE GAP The evidence base is laboratory and field-experiment work plus the NIST taxonomy. A paper implying that real-world defence adversarial attacks are documented would be overstating the record, and I would rather say so than imply otherwise.

The taxonomy to anchor policy language is NIST AI 100-2e2025, published in March 2025 and covering adversarial machine learning attacks and mitigations, with a corrected version issued on 1 April 2025 and an error notice on 3 June 2025.[4] STANDARD It is a United States publication with no UK legal force and it is the best available common vocabulary, which matters more than it sounds when a supplier and a customer are arguing about whether a threat model is adequate.

6.2 Classified and controlled information, and the export control that reaches data

The classification problem in this sector is not primarily about the model. It is about what crosses a boundary when the model is used, since prompts, telemetry, drawings and outputs can each carry controlled content across a security or export boundary without anybody intending it.

UK national controls on AI-relevant items exist as entries in the strategic export control lists and not as a control on AI as such. Three national entries added in April 2024 cover semiconductor technologies and equipment, quantum computing technologies and additive manufacturing, with the semiconductor entry covering high-end digital processing integrated circuits most suitable for AI training.[17] STANDARD Amending regulations made on 28 April 2025 and in force from 20 May 2025 made administrative amendments to the technical parameters of that entry and updated the military list.

There is no UK export control on model weights as such. The control bites on compute hardware and on technology and software transfer, and JSP 936 makes the practical point for a defence programme, which is that assessment of restrictions caused by foreign export controls should include data as well as algorithms and models, so that MOD retains freedom of action to maintain, modify, upgrade and operate the AI.[5] STANDARD

On the United States side two changes matter for UK companies and both cut against received wisdom. The ITAR exemption for AUKUS took effect on 1 September 2024, creating a new exemption and an excluded technologies list, and permitting licence-free transfers between authorised parties in the three nations subject to conditions, with a final rule published on 30 December 2025 and effective the same day.[35] STANDARD And the framework for artificial intelligence diffusion, published on 15 January 2025 with a compliance date of 15 May 2025, was rescinded on 13 May 2025, two days before it would have applied.[36] STANDARD A paper describing the diffusion rule as operative is describing something that never took effect.

6.3 Software provenance, where the mandate went backwards

Anybody writing a supply chain policy in 2026 needs to know that the United States federal software attestation regime was rolled back. Office of Management and Budget memorandum M-26-05, issued on 23 January 2026, rescinds the two memoranda that had required agencies to obtain self-attestation of conformance to the secure software development framework before use. Use of the common attestation form becomes optional, and agencies must instead develop assurance policies matching their own risk determinations and mission needs, and may still require software bills of materials.[37] STANDARD

So United States federal contracts do not mandate software bills of materials or attestation as at July 2026. The position is agency by agency and risk-based. A supplier that has built its provenance capability anyway is in a good position, since the requirement will return through contract terms, and it is a great deal easier to have the capability and not need it.

Two requirements that are live and dated. The cybersecurity maturity model certification rule was published in the Federal Register on 10 September 2025 and took effect on 10 November 2025, with a first phase running to 9 November 2028 in which the clause applies where a programme manager or requiring activity determines to apply it, and a second phase from 10 November 2028 in which it applies wherever contractors will handle federal contract information or controlled unclassified information on their own systems.[38] STANDARD Any UK entity in a United States defence supply chain handling that material is inside this by November 2028 at the latest and can be pulled in earlier.

Three cross-cutting references sit underneath all of this and are worth stating once. The National Cyber Security Centre's guidelines for secure AI system development of November 2023 give the lifecycle structure most customers will expect a supplier to have adopted. The NIST generative AI profile of July 2024 supplies a risk taxonomy, and ISO/IEC 42001 of December 2023 supplies a certifiable management system, neither of which carries UK legal force and both of which appear in supplier assurance questionnaires. Where a programme touches personal data, the Information Commissioner's accountability expectations apply in the ordinary way.[1,2,3,50] STANDARD

On the UK side the Defence Cyber Protection Partnership operates through the Cyber Security Model, DEFCON 658 and Def Stan 05-138, with a cyber risk assessment on the contract producing a risk profile, which drives a supplier assurance questionnaire and flows down to sub-contracts.[39] STANDARD Cyber Essentials is the baseline at the lower profiles.

NATO adds principles and not a certification scheme. The revised Artificial Intelligence Strategy released on 10 July 2024 retains the six principles of responsible use, being lawfulness, responsibility and accountability, explainability and traceability, reliability, governability, and bias mitigation.[40] STANDARD Suppliers are expected to develop systems in accordance with Alliance requirements and those principles. There is no NATO AI certification standard for industry, and a supplier being told there is should ask which document.

CHAPTER SEVEN

Human factors and the accountability gap

This is the oldest and best-developed evidence base in the whole paper, and it is the one AI programmes most consistently ignore.

Aviation has been studying what automation does to the people supervising it for more than 40 years, and the findings are stable, replicated and inconvenient. Any AI programme in this sector that treats human factors as a user experience activity is discarding the most reliable evidence available to it.

The founding text is Bainbridge's paper on the ironies of automation, published in 1983, whose argument is that automating the easy parts of a task leaves the human with the hardest residual work while degrading the skill needed to do it.[41] PEER-REVIEWED Sarter and Woods established mode error and automation surprise in supervisory control, and documented operational experience of automation surprises on a specific aircraft type.[42] PEER-REVIEWED Parasuraman and Riley set out the use, misuse, disuse and abuse taxonomy, and Parasuraman and Manzey later synthesised the evidence on complacency and automation bias as an attentional problem.[43] PEER-REVIEWED

The regulatory apparatus has absorbed this. The flight deck automation working group's final report of autumn 2013 produced 28 interconnected data-driven findings and 18 recommendations, on themes of complexity in systems and operations, concerns about degradation of pilot knowledge and skills, and integration and interdependence, noting that several accidents had occurred where investigative reports identified vulnerabilities similar to those the report identified.[44] REGULATOR

7.1 Two quantified studies worth knowing by their numbers

A field survey of 200 Dutch airline pilots, asked about their most recent automation surprise, found roughly three such events per pilot per year, with a reportable incident estimated at once every three years per pilot. Unchanged trust afterwards was reported by 59 per cent, and only 9 per cent substantially reduced their trust. Detection was overwhelmingly self-detection at 89 per cent, with 7 per cent via alerting systems and 5 per cent via the other pilot. Attributed causes were system misunderstanding at 63 per cent, incorrect manual input at 24 per cent, knowledge gaps at 19 per cent, excessive trust or complacency at 12 per cent, and loss of situational awareness at 8 per cent.[45] PEER-REVIEWED

The cleanest modern experimental demonstration comes from a full flight simulator study of 20 professional pilots flying six landing scenarios at three automation levels. Higher automation increased flight performance and reduced mental workload, and was associated with a decrease in vigilance to primary instruments, particularly flight path indicators and engine thrust, which the authors read as confirming the risks of adverse effects of automation on visual monitoring.[46] PEER-REVIEWED

THE TRADE, STATED EXPERIMENTALLY

Automation buys performance and reduces workload. It costs monitoring. Both halves are measured in the same study.

That result is the one to put in front of a programme board proposing to introduce a decision aid into an operational crew position. The performance improvement is real and it is not free, and the cost is paid in an attentional resource that nobody is currently measuring.

7.2 Two cases, described accurately

The 737 MAX accidents. The Indonesian investigation into the Lion Air accident of 29 October 2018 published its final report on 25 October 2019, listing nine contributing factors presented chronologically and explicitly not ranked by degree of contribution. They include incorrect assumptions about pilot response during design and certification, reliance on those assumptions plus an incomplete system review leading to single-sensor certification approval, a design vulnerability arising from dependence on a single angle-of-attack sensor, insufficient guidance in flight manuals and crew training, an alert not being enabled as intended, a miscalibrated replacement sensor installed undetected, installation testing that failed to identify the miscalibration, missing documentation of the prior flight's problems, and multiple alerts and repeated activations exceeding crew management capacity.[47] CASE LAW

The Ethiopian investigation into the accident of 10 March 2019 published its final report dated 23 December 2022, giving the cause as repetitive nose-down trim command activation due to erroneous angle-of-attack sensor input which the flight crew was unable to overcome, with contributing factors including inadequate design and certification that failed to account for single-sensor failure, insufficient pilot training, unclear runaway stabiliser procedures, and deficiencies in system safety assessment and regulatory oversight of certification delegation.[48] CASE LAW That report is formally contested. The United States National Transportation Safety Board issued comments and a press release on 27 December 2022 objecting that its comments were not appended to the final report as the international convention requires, and it disputes elements of the technical analysis.[48] CASE LAW

Now the discipline. The system at the centre of these accidents was not artificial intelligence and it was not machine learning. It was a deterministic control law fed by a single sensor. Its relevance here is entirely about automation authority, single-point sensor dependence, undisclosed system behaviour, assumed crew response times and certification delegation, every one of which recurs in AI assurance. Calling it an AI accident is wrong and a specialist reader will stop reading.

A machine learning perception failure with an official finding. The National Transportation Safety Board's report on the collision between a vehicle controlled by a developmental automated driving system and a pedestrian in Tempe, Arizona in March 2018 was adopted on 19 November 2019. The probable cause given is the failure of the vehicle operator to monitor the driving environment and the operation of the automated driving system because she was visually distracted by her personal telephone, with contributing factors including the developer's inadequate safety risk assessment, ineffective operator oversight, insufficient mechanisms for addressing automation complacency stemming from poor safety culture, the pedestrian's impaired midblock crossing, and insufficient state oversight of automated vehicle testing.[49] CASE LAW

The machine learning finding inside it is the one to carry into an assurance case. The system first detected the pedestrian 5.6 seconds before impact and never accurately classified her as a pedestrian or predicted her path, changing the classification several times between vehicle, bicycle and other, which prevented consistent path prediction. Only 1.2 seconds before impact did it identify an emergency, by which point avoidance would have required braking beyond its design specifications.[49] CASE LAW

Read that carefully because it is not the failure mode people expect. The classifier was not blind. It detected the object at 5.6 seconds, which is a long time. What failed was classification stability, and unstable classification destroyed track continuity, which destroyed path prediction. Detection performance metrics would have looked acceptable. This is directly the accountability gap that JSP 936 requires analysis of at paragraph 125(b), and directly the human factors building block in the European roadmap, in a real accident with an official finding.[5,9] STANDARD

7.3 What to require, given all of that

Four requirements follow, and I would write all four into a specification for anything placed in front of an operator.

HUMAN FACTORS REQUIREMENTS FOR AN OPERATIONAL DECISION AID

01 **Stability reported**	02 **Authority explicit**	03 **Monitoring measured**	04 **Degradation rehearsed**
The system reports classification or recommendation stability over time, and not only its confidence at a point. Instability is displayed as a state, because it is one.	The operator's authority to override is stated on the interface, tested in trials, and evidenced by overrides actually occurring in service.	Visual attention or an equivalent proxy is measured during human factors trials, because the published evidence says monitoring is what automation costs.	Crews practise the task without the aid on a defined cycle, and the proportion who have done so recently is a reported figure.

FIGURE 1 Four requirements drawn from the human factors literature.

Proposed specification content. ECAVEO The monitoring requirement follows the full flight simulator finding on vigilance to primary instruments,[46] and the stability requirement follows the unstable-classification mechanism in the Tempe investigation.[49]

The fourth is the one that gets cut on cost, and it is the one the accountability gap turns on. An operator who has not flown, sailed or run the task manually within a year is not exercising meaningful control over the automation, whatever the interface says. PRACTICE

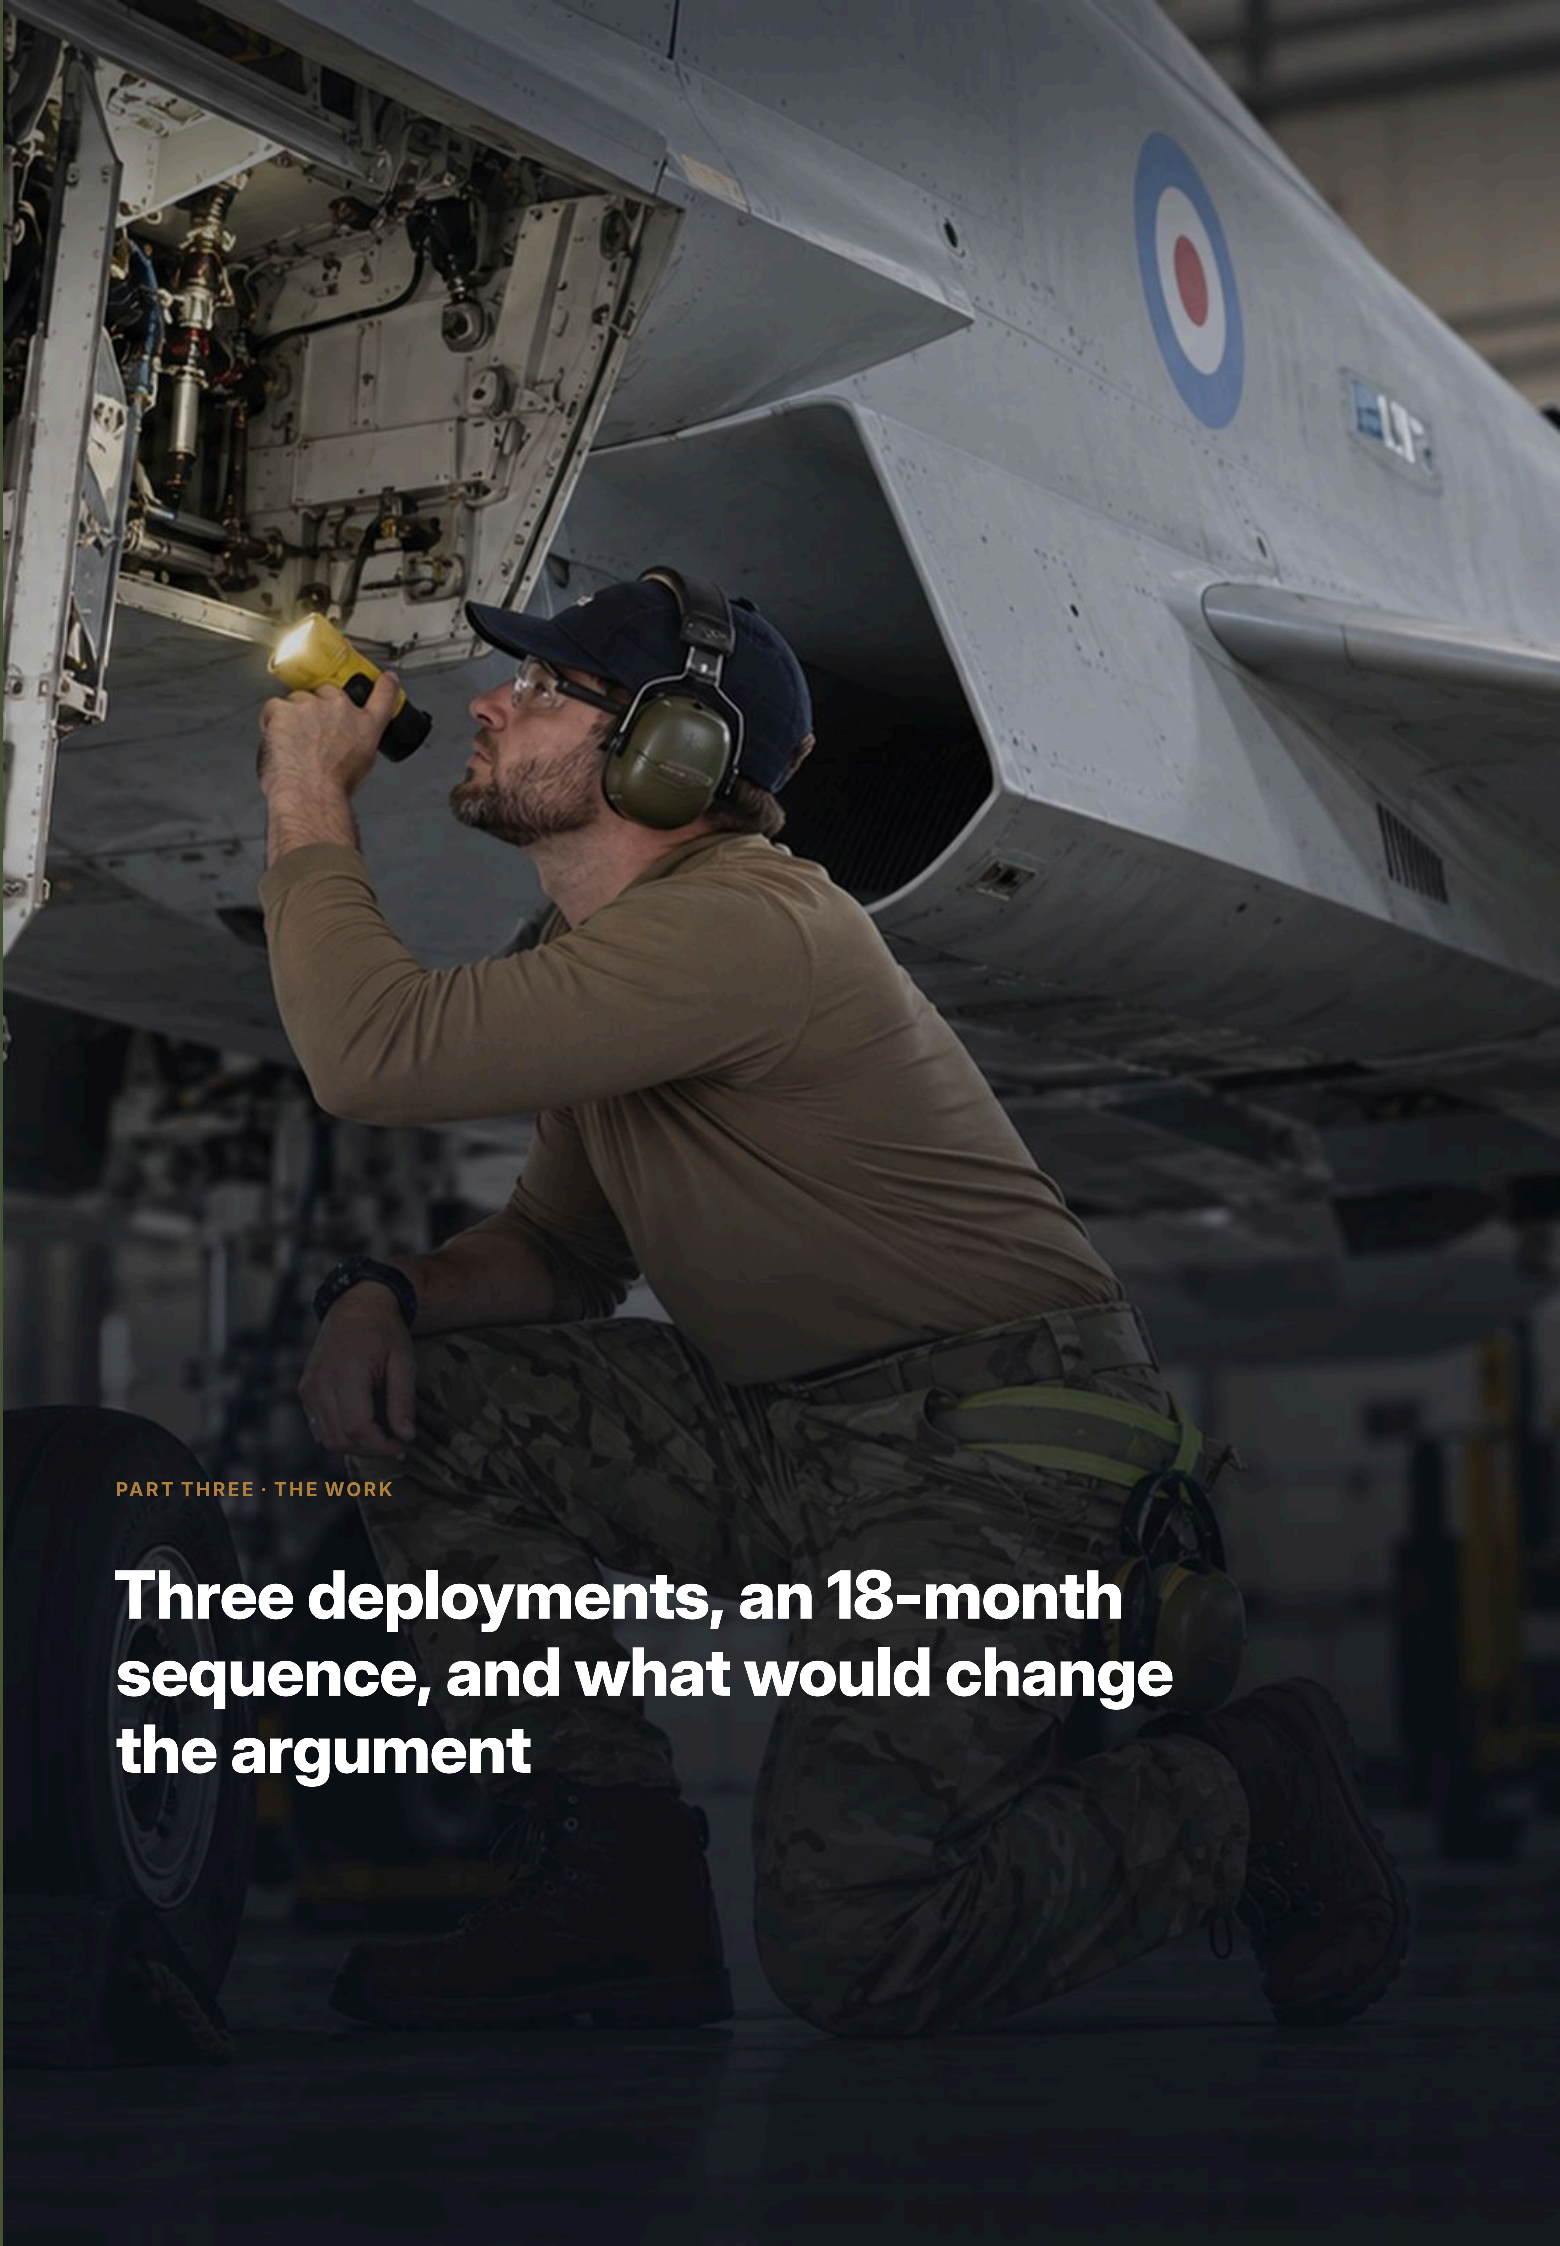

PART THREE · THE WORK

Three deployments, an 18-month sequence, and what would change the argument

CHAPTER EIGHT

Three deployments, and how to evaluate them

Each is bounded, evaluable inside a quarter, and produces evidence in a form the assurance case can absorb. None of them flies.

8.1 The maintenance evidence assistant

Scope. Retrieve approved manuals and configuration-matched records for one platform and one task. The system returns source-linked extracts. It writes nothing to the maintenance record, raises no work order and makes no airworthiness determination.

Why this one. Matching a maintainer to the right technical information for the actual configuration in front of them is a retrieval problem, and it is one the industry already fails at expensively. The evidence produced feeds the assurance case directly, because citation accuracy and configuration match are exactly the properties a regulator will ask about.

Evidence to collect. Citation accuracy, meaning whether every returned extract exists and says what the system reports. Configuration match rate against the recorded build state. Abstention rate, meaning how often the system declines when it has no confident match, which should be non-zero. Maintainer time to correct information. Unsafe answer rate, defined before the pilot with the maintenance authority and counted separately from ordinary errors.

Stop condition. Any unsafe answer that a maintainer did not catch. A single instance ends the pilot and triggers a review, because the whole value proposition rests on the assumption that the maintainer catches them.

8.2 The requirements traceability reviewer

Scope. Identify gaps and conflicts across one controlled requirement and verification set. Output is a list of candidate findings for an engineer to disposition. Nothing is closed and no requirement is edited.

Why this one. It sits inside the engineering record where configuration control already exists, so the pilot generates auditable evidence without new infrastructure. The industrial validation gap in the literature means anybody who runs this properly will know more than the published field does.[27] PEER-REVIEWED

Evidence to collect. Recall against a held-out set of known defects the system has not seen, which requires that set to be built first and is the real work of the pilot. False positive rate at the working threshold. Engineer effort per finding dispositioned. Audit traceability, meaning whether the reasoning behind each candidate finding survives into the record.

Stop condition. A false positive rate that engineers stop working through. Set the threshold with the engineers before starting, and write it down, because it will be argued about afterwards.

8.3 The factory quality adviser

Scope. A vision model in shadow mode on one inspection point. Inspectors retain disposition entirely and do not see the model's output during the shadow period.

Why this one. Section 4.3 sets out the benchmark trap, and shadow mode on one real line is the direct answer to it. It also produces the cross-dataset evidence the published literature lacks, on the only dataset that matters to the company.

Evidence to collect. Detection performance by defect type, since aggregate accuracy conceals the rare classes that matter. False reject rate, which is what determines whether production will tolerate it. Drift across the shadow period, particularly across shift changes and part mix changes. Throughput effect. Override analysis once live, including which inspectors override and on what.

Stop condition. Any defect class where detection is materially worse than the incumbent inspection, unless that class is explicitly excluded from scope and the exclusion is recorded.

ILLUSTRATIVE SCENARIO

Ashcombe Systems runs the maintenance retrieval pilot

An invented tier one supplier runs the retrieval assistant across one rotorcraft variant and one task family for 10 weeks. Citation accuracy comes out at 97 per cent. The configuration match rate is 82 per cent, which is worse than expected and turns out to be the finding.

Investigation shows the shortfall is not a model problem. The build state records for 11 airframes had been updated in the maintenance system and not in the technical publication index, so the system was correctly retrieving against a record that was itself wrong. Nobody had noticed, because maintainers had been resolving the discrepancy from memory for years.

The pilot's most valuable output is therefore a data quality defect that predates the AI by a decade, and a business case argument the programme did not expect to make, which is that the retrieval capability is worth building partly because it makes the underlying record errors visible.

CHAPTER NINE

From demonstrator to fleet

An 18-month sequence that produces assurance evidence as a by-product of engineering, so that the evidence exists when a contract asks for it.

9.1 The first six weeks

Select the mission or engineering problems, and define the security, safety, configuration and authority boundaries for each. Boundary definition before capability selection is the ordering that matters, and it is the one most programmes invert.

Write the operational design domain for each candidate, in a form that can be tested. If it cannot be written down, that is the finding, and it means the capability is not ready for a programme.

Characterise the baseline. Where the substitution argument in section 1.2 will be used, measure the incumbent, whether human or legacy system, on the same input distribution. Budget for this separately and expect it to take most of the six weeks.

Establish the evidence chain. Data provenance from collection to ingestion, configuration management for all output-influencing data including metadata, and a versioned record of model, data, tooling and test.[5] STANDARD Retrofitting this is the single largest avoidable cost in the sector.

Identify the export control and classification boundaries the work will cross, covering data as well as algorithms and models.

9.2 The rest of the sequence

THE 18-MONTH SEQUENCE

PHASE	TIMING	OUTPUTS
Bound	Weeks 1 to 6	Problem selection, operational design domain written and testable, baseline characterised, evidence chain established, security and export boundaries mapped.
Build	Months 2 to 5	Three bounded pilots under configuration control, none of them authoritative and one in shadow mode, with held-out evaluation sets, a threat model and adversarial testing, and human factors trials in which monitoring is measured.
Verify	Months 6 to 11	Independent verification and validation across the domain including boundary conditions and edge cases, demonstrated safe behaviour outside the domain, regression suite established, release authority defined.
Field	Months 12 to 18	Controlled introduction with monitoring, drift detection, change control tied to a compatibility matrix, degraded-mode rehearsal on a defined cycle, and through-life evidence retention.

TABLE 4 Each phase produces evidence the next depends on. The independent verification in phase three is what converts a demonstrator into something a customer can accept.

Two items in that table are the ones that separate an assurance case from a slide deck. Demonstrated safe behaviour outside the operational design domain is a specific requirement of the customer's directive and it is routinely interpreted as a statement that the system will not be used outside the domain, which is a different and much weaker claim.[5] STANDARD And a compatibility matrix tied to change control is what stops a model update, a data update, a hardware refresh or a platform modification silently invalidating evidence that was correct when it was produced.

9.3 What good looks like after 18 months

A fielded capability with measured performance against the baseline set in week one. An assurance case that a third party can follow from need through configuration and test into service, without needing to interview anybody. A defined release authority, exercised at least once. A regression suite that has caught something. Crews or maintainers who have operated the task without the aid within the last year, recorded as a proportion. And a backlog that separates proven capability from research and from things the company has decided the evidence will not support.

That last category is the one a customer will find most reassuring, and the one commercial pressure most reliably erases.

CHAPTER TEN

Where the evidence is thin

The claims in this paper that would change if the evidence changed, and the places where I could not get to the primary source.

10.1 What I could not establish

No measured, independently verified improvement in UK military aircraft or vehicle availability attributable to AI-based predictive maintenance was found. The published economic case is simulation and life-cycle cost modelling.[23,24,31] EVIDENCE GAP

No verified real-world adversarial machine learning attack on a fielded defence system was found. The evidence base is laboratory and field-experiment work.[33,34] EVIDENCE GAP

No recent aviation accident or serious incident with a machine learning system as a causal factor was found. The nearest official finding involving a machine learning perception system is the road vehicle investigation used in Chapter 7.[49] EVIDENCE GAP

Sample sizes and effect sizes for the aerospace requirements question-answering study could not be obtained, and I have cited the paper without relying on its results.[28] EVIDENCE GAP The same caution applies to the surface defect systematic review, whose existence and venue I verified without obtaining its quantitative findings.[29] EVIDENCE GAP

The full text of Def Stan 00-055 issue 5 and Def Stan 00-056 issue 8 sits behind MOD ASSIST and was not inspected. Issue numbers and dates are verified from standards records.[15] EVIDENCE GAP Similarly, the precise publication date of Def Stan 05-138 issue 4 was not confirmed.[39] EVIDENCE GAP

Fleet-wide serviceability percentages for the Royal Air Force are not routinely published, and anybody quoting one should be asked for the source. EVIDENCE GAP

10.2 What would change the argument

Publication of ED-324 and ARP6983. The whole of Chapter 3 is written against a draft, and the published Issue 1 would replace judgement with a process a company can follow.[12] STANDARD

Publication of JSP 936 Part 2, which would set out the Responsible AI Senior Officer's delegated responsibilities and the governance chain, and which would change what a supplier is engaging with.[5] STANDARD

Any well-designed field evaluation of predictive maintenance with a characterised baseline. One such study would settle Chapter 5, in either direction, and I would revise the chapter on the strength of it.

An update to the Civil Aviation Authority strategy documents. The two strategy documents have stood at the same version since December 2024 while the European position moved substantially, and a UK rulemaking position would change how a dual-market company should sequence its work.[13] STANDARD

A published MAA position on AI or machine learning assurance. Its absence is currently a planning assumption in Chapter 3, and one I would rather not be relying on.[14] EVIDENCE GAP

10.3 Where I have gone beyond the sources

The reading of the F-35 availability figures in Chapter 5 is mine. The Government Accountability Office attributes the decline to parts, depot capacity, the reset programme and contractual incentives, and says nothing about analytics.[22] REGULATOR My inference is that a fleet this instrumented is the best available natural test of the proposition that data richness produces availability, and that the test has not gone the way the market assumes. A reasonable person could argue that the confound is too large to read anything from. PRACTICE

The recommendation to isolate learned components behind a machine learning constituent boundary now, ahead of ED-324 publication, is mine. It follows the draft's architecture and the draft may change. ECAVEO

The four human factors requirements in Chapter 7 are my synthesis of a literature that does not itself express them as requirements. ECAVEO

And there is one thing I have turned over repeatedly without reaching a settled view. JSP 936 says that in some applications assurance risk will simply present too much risk and AI must not be adopted.[5] STANDARD That is the right principle. What I cannot work out is who, in practice, is positioned to invoke it. The supplier has a commercial interest against it. The programme office has a schedule. The Responsible AI Senior Officer is at top-level budget level and there are 17 of them across the whole of Defence, which is not a ratio that supports case-by-case scrutiny.[8] REGULATOR The Ministry's AI Ethics Advisory Panel, established in 2021, is explicitly advisory and holds no decision-making authority and is not a formal board within the governance structures.[7] REGULATOR I suspect the answer is the chief engineer, on the strength of professional obligation rather than institutional position, which is a thin thread to hang a prohibition on. I have not found anybody with a better answer, and I would be glad to be told one.

References

SOURCES, WITH METHOD, STATUS AND PROVENANCE WHERE IT BEARS ON THE WEIGHT A CLAIM CAN CARRY

Important claims were checked against government departments, regulators, standards bodies, accident investigation authorities and primary official material. Where a standard is in draft, the reference says so and gives its status at the date of writing. Where a source is research from a body with a commercial or membership interest in the result, the reference says so.

1. Information Commissioner's Office. *Guidance on AI and data protection*. Last updated 15 March 2023, carrying a notice that it is under review following the Data (Use and Access) Act. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/
2. National Cyber Security Centre. *Guidelines for secure AI system development*. Version 1.0, 27 November 2023. Endorsed and co-sealed by agencies from 17 other countries, and developed with 21 other international agencies and ministries. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development
3. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, released 26 July 2024, with no subsequent revision. A United States voluntary framework with no UK legal force. https://doi.org/10.6028/NIST.AI.600-1
4. National Institute of Standards and Technology. *Adversarial Machine Learning: a taxonomy and terminology of attacks and mitigations*, NIST AI 100-2e2025, March 2025, announced 24 March 2025, corrected PDF 1 April 2025, error notice 3 June 2025. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2025.pdf
5. Ministry of Defence. *JSP 936: Dependable Artificial Intelligence (AI) in Defence, Part 1: Directive*. Version 1.1, November 2024, first published 13 November 2024, 41 pages. Part 2 is referenced at paragraph 100 and has not been published. Paragraph references in the text are to this version. https://www.gov.uk/government/publications/jsp-936-dependable-artificial-intelligence-ai-in-defence-part-1-directive
6. Ministry of Defence. *Putting artificial intelligence at the heart of UK Defence*. 10 June 2026. Published as correspondence, a letter to Defence personnel, and not as a strategy or policy instrument. It contains no funding figures. https://www.gov.uk/government/publications/putting-artificial-intelligence-ai-at-the-heart-of-uk-defence
7. Ministry of Defence. *Ambitious, safe, responsible: our approach to the delivery of AI-enabled capability in Defence*. 15 June 2022. Sets the five ethical principles, being Human-Centricity, Responsibility, Understanding, Bias and Harm Mitigation, and Reliability. Published alongside the *Defence Artificial Intelligence Strategy* of the same date. See also the *Ministry of Defence AI Ethics Advisory Panel*, established 2021, an advisory body without decision-making authority. https://www.gov.uk/government/groups/ministry-of-defence-ai-ethics-advisory-panel https://www.gov.uk/government/publications/ambitious-safe-responsible-our-approach-to-the-delivery-of-ai-enabled-capability-in-defence
8. Ministry of Defence. *Laying the Groundwork: Responsible AI Senior Officers' Report 2025*. Corporate report, 3 October 2025, 12 pages. Records 17 nominated officers. A qualitative maturity report containing no quantitative programme data. https://www.gov.uk/government/publications/laying-the-groundwork-responsible-ai-senior-officers-report-2025
9. European Union Aviation Safety Agency. *Artificial Intelligence Roadmap 2.0: a human-centric approach to AI in aviation*. Version 2.0, May 2023, announced 10 May 2023. No version 3.0 exists. https://www.easa.europa.eu/en/newsroom-and-events/news/easa-artificial-intelligence-roadmap-20-published
10. European Union Aviation Safety Agency. *Concept Paper: guidance for safety-related artificial intelligence applications*, proposed Issue 03, June 2026, released for a 10-week comment period on 3 June 2026; and Issue 02, *guidance for Level 1 and 2 machine learning applications*, 6 March 2024. https://www.easa.europa.eu/sites/default/files/dfu/easa_concept_paper_guidance_for_artificial_intelligence_applications_proposed_issue_03.pdf
11. European Union Aviation Safety Agency. *NPA 2025-07*, the Agency's first regulatory proposal on artificial intelligence for aviation, opened for consultation 10 November 2025 for three months. A proposal, not a rule. https://www.easa.europa.eu/en/newsroom-and-events/news/easas-first-regulatory-proposal-artificial-intelligence-aviation-now-open
12. EUROCAE Working Group 114 and SAE G-34. *ED-324 / ARP6983, Development and assurance considerations for aeronautical systems and equipment implemented with machine learning*, also described as a process standard for development and certification or approval of aeronautical safety-related products implementing AI. Draft. EUROCAE lists a target publication date of 31 December 2026. A joint status presentation of August 2025 recorded draft seven with approximately 1,800 first-ballot comments. Issue 1 scope is non-

adaptive supervised machine learning up to DAL C, AL 3 and SWAL 2, excluding information security and human factors. https://www.eurocae.net/working-group/wg-114/

13. Civil Aviation Authority. *CAP3064, The CAA's response to emerging AI-enabled automation*, and *CAP3064A, Part A: strategy for regulating AI and advanced automation in aerospace*, both version 1.3.2, 3 December 2024; and *CAP3019, AI in aviation: a technology outlook*, version 1.1, 12 August 2024. Both remained at these versions as at July 2026, although the CAP3064 series has since gained further parts, including CAP3064G on consumer and AI principles, version 1, 11 April 2025. https://www.caa.co.uk/about-us/research-and-innovation/artificial-intelligence/

14. Military Aviation Authority. *MAA Regulatory Publications*, GOV.UK collection, published 11 November 2014; and *Regulatory Article 1600, Uncrewed Air Systems Categorization*, Issue 10, page last updated 28 November 2025. No dedicated MAA regulatory article, manual or position paper on AI or machine learning assurance was located. https://www.gov.uk/government/collections/maa-regulatory-publications

15. Ministry of Defence. *Def Stan 00-055 Part 1, Requirements for safety of programmable elements in defence systems*, Issue 5, 28 July 2021; and *Def Stan 00-056 Part 1, Safety management requirements for defence systems*, Issue 8, 14 October 2023. Distributed through MOD ASSIST; issue numbers and dates verified from standards records, full text not inspected for this paper.

16. Ministry of Defence. *Defence Artificial Intelligence Playbook*. GOV.UK first published 1 February 2024; the document cover is dated January 2024. Contains eight case studies and no measured outcomes. https://www.gov.uk/government/publications/defence-artificial-intelligence-ai-playbook

17. Department for Business and Trade, Export Control Joint Unit. *UK strategic export control lists: the consolidated list*, last updated 16 December 2025, including national entries PL9013, PL9014 and PL9015 added in April 2024; and *The Export Control (Amendment) Regulations 2025*, SI 2025/532, made 28 April 2025, laid 29 April 2025, in force 20 May 2025. https://www.legislation.gov.uk/uksi/2025/532/contents/made

18. Ministry of Defence. *MOD Trade, Industry and Contracts 2025*. Published 25 September 2025, revised 12 March 2026, covering financial year 2024/25. https://www.gov.uk/government/statistics/mod-trade-industry-and-contracts-2025

19. House of Commons Library. *UK defence expenditure*, Research Briefing CBP-8175, 22 July 2026; and Ministry of Defence, *Defence Departmental Resources 2025*, published 15 January 2026. https://commonslibrary.parliament.uk/research-briefings/cbp-8175/

20. ADS Group. *Industry Facts and Figures 2026*, data year 2025, published July 2026. Trade association figures compiled from member data. The aerospace-only split is published in the full report at pages 8 and 9. The defence export figure is a five-year average and the gross value added figure is an estimate for 2025. https://www.adsgroup.org.uk/facts2025/

21. Defence Committee. *Developing AI capacity and expertise in UK defence*. Second Report of Session 2024 to 2025, 10 January 2025. https://publications.parliament.uk/pa/cm5901/cmselect/cmdfence/590/report.html

22. United States Government Accountability Office. *F-35 Sustainment: actions needed to ensure updated strategy improves persistent readiness challenges*, GAO-26-108113, 11 June 2026. A United States report on a United States programme; comparable UK figures are not published. https://www.gao.gov/products/gao-26-108113

23. Hölzel, N. B. and Gollnick, V. *Cost-benefit analysis of prognostics and condition-based maintenance concepts for commercial aircraft considering prognostic errors*. Annual Conference of the PHM Society, 7(1), 2015. A discrete-event simulation study, not a measured operational result.

24. Scanff, E., Feldman, K. L., Ghelam, S., Sandborn, P., Glade, M. and Foucher, B. *Life cycle cost impact of using prognostic health management for helicopter avionics*. Microelectronics Reliability, 47(12), 2007, pp. 1857 to 1864. Life-cycle cost modelling, not a measured fleet outcome.

25. Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T. *The effects of generative AI on high-skilled work: evidence from three field experiments with software developers*. Management Science, 27 February 2026, DOI 10.1287/mnsc.2025.00535. Three randomised trials, 4,867 developers, at Microsoft, Accenture and a third large firm. The tool studied is a Microsoft product and two of the three sites are the vendor and a systems integrator.

26. Becker, J., Rush, N., Barnes, E. and Rein, D. *Measuring the impact of early-2025 AI on experienced open-source developer productivity*. arXiv:2507.09089, 12 July 2025, revised 25 July 2025. Randomised, 16 developers, 246 tasks. Very small sample of self-selected experts on mature codebases; does not generalise. https://arxiv.org/abs/2507.09089

27. Zadenoori, M. A., Dąbrowski, J., Alhoshan, W., Zhao, L. and Ferrari, A. *Large language models for requirements engineering: a systematic literature review*. arXiv:2509.11446, 14 September 2025. 74 primary studies from 2023 and 2024. The review records an industrial validation gap. https://arxiv.org/abs/2509.11446

28. Yang, L., Luo, Y., Gao, H., and others. *Evaluating large language models for requirements question answering in industrial aerospace software*. Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering, 23 June 2025, pp. 366 to 377. Cited for existence and relevance; the full text was not obtainable and no result from it is relied on here.

29. Ameri, R., Hsu, C.-C. and Band, S. S. *A systematic review of deep learning approaches for surface defect detection in industrial applications*. Engineering Applications of Artificial Intelligence, 130, article 107717, April 2024. Existence and venue verified; quantitative findings were not obtainable and none is attributed to it here.

30. da Silva, Tavares, Lopes and Silva. *Cross-dataset benchmarking of deep learning models for surface defect classification in metal parts*. Applied Sciences, 16(6), article 3022, 20 March 2026. Cross-dataset testing is where generalisation failure becomes visible. https://www.mdpi.com/2076-3417/16/6/3022

31. Stanton, I., Munir, K., Ikram, A. and El-Bakry, M. *Predictive maintenance analytics and implementation for aircraft: challenges and opportunities*. Systems Engineering, 26(2), 2023, pp. 216 to 237. A systematic literature review, which records that no comprehensive survey devoted to the aircraft manufacturing industry existed and that research focus is biased towards engines for want of public datasets.

32. Mutlu, U. and Kaewunruen, S. *Digitalised predictive maintenance in railways: a systematic review of AI, BIM and digital twins*. Infrastructures, 11(3), article 87, 8 March 2026. Cited here for the structural finding that safety-critical infrastructure replaces components before failure, producing a shortage of run-to-failure training data. https://www.mdpi.com/2412-3811/11/3/87

33. Eykholt, K., and others. *Robust physical-world attacks on deep learning visual classification*. Computer Vision and Pattern Recognition, 2018; arXiv:1707.08945. https://arxiv.org/abs/1707.08945

34. Chu, S. and Chen, Y. *On the adversarial robustness of aerial detection*. Frontiers in Computer Science, 6, article 1349206, 2024. A survey covering optical, infrared and synthetic aperture radar sensing. Its own stated limitations include insufficient real-world testing, no standardised defence evaluation framework, and reliance on low-resolution datasets. https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2024.1349206/full

35. United States Department of State, Directorate of Defense Trade Controls. ITAR exemption for defence trade and cooperation among Australia, the United Kingdom and the United States. Interim final rule published in the Federal Register 20 August 2024, effective 1 September 2024, creating ITAR section 126.7 and the Excluded Technologies List; final rule published in the Federal Register 30 December 2025, effective the same day. https://www.state.gov/key-elements-of-the-international-traffic-in-arms-regulations-exemption-for-defense-trade-and-cooperation-among-australia-the-united-kingdom-and-the-united-states

36. United States Bureau of Industry and Security. *Framework for Artificial Intelligence Diffusion*, published 15 January 2025 with a compliance date of 15 May 2025, rescinded by announcement of 13 May 2025 and replaced by three guidance documents. It never took effect. https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens

37. United States Office of Management and Budget. *M-26-05, Adopting a risk-based approach to software and hardware security*, 23 January 2026, rescinding M-22-18 and M-23-16. Self-attestation via the common form becomes optional and assurance becomes agency-specific and risk-based.

38. United States Department of Defense. *Defense Federal Acquisition Regulation Supplement: assessing contractor implementation of cybersecurity requirements*, DFARS Case 2019-D041. Federal Register 10 September 2025, effective 10 November 2025. Phase 1 to 9 November 2028; Phase 2 from 10 November 2028. https://www.federalregister.gov/documents/2025/09/10/2025-17359/defense-federal-acquisition-regulation-supplement-assessing-contractor-implementation-of

39. Ministry of Defence. *Defence Cyber Protection Partnership*. GOV.UK guidance published 12 September 2019, last updated 2 January 2025. Operates through the Cyber Security Model, DEFCON 658 and Def Stan 05-138. The publication date of Def Stan 05-138 Issue 4 was not confirmed for this paper. https://www.gov.uk/guidance/defence-cyber-protection-partnership

40. North Atlantic Treaty Organization. *Summary of the NATO Artificial Intelligence Strategy*, first adopted October 2021, revised strategy released 10 July 2024. Sets six principles of responsible use. There is no NATO AI certification standard for industry. https://www.nato.int/cps/en/natohq/official_texts_227237.htm

41. Bainbridge, L. *Ironies of automation*. Automatica, 19(6), 1983, pp. 775 to 779.

42. Sarter, N. B. and Woods, D. D. *How in the world did we ever get into that mode? Mode error and awareness in supervisory control*. Human Factors, 37(1), 1995, pp. 5 to 19; and *Team play with a powerful and independent agent: operational experiences and automation surprises on the Airbus A-320*. Human Factors, 39(4), 1997.

43. Parasuraman, R. and Riley, V. *Humans and automation: use, misuse, disuse, abuse*. Human Factors, 39(2), 1997, pp. 230 to 253; and Parasuraman, R. and Manzey, D. H. *Complacency and bias in human use of automation: an attentional integration*. Human Factors, 52(3), 2010, pp. 381 to 410.

44. Federal Aviation Administration, PARC and CAST Flight Deck Automation Working Group. *Operational use of flight path management systems*. Final report, autumn 2013. 28 findings and 18 recommendations. https://skybrary.aero/bookshelf/books/2501.pdf

45. de Boer, R. and Dekker, S. *Models of automation surprise: results of a field survey in aviation*. Safety, 3(3), article 20, 2017. Questionnaire study of 200 Dutch airline pilots recruited through the pilots' association, pilot websites and pre-flight briefings. No funder stated.

46. Causse, M., Mercier, M., Lefrançois, O. and Matton, N. *Impact of automation level on airline pilots' flying performance and visual scanning strategies: a full flight simulator study*. Applied Ergonomics, 125, article 104456, 2025. 20 professional pilots, six landing scenarios, three automation levels.

47. Komite Nasional Keselamatan Transportasi. Final report on the accident to PK-LQP, Lion Air flight JT610, 29 October 2018. Published 25 October 2019. Nine contributing factors, presented chronologically and expressly not ranked by degree of contribution.

48. Ethiopian Aircraft Accident Investigation Bureau. Final report AI-01/19 on the accident to ET-AVJ, Ethiopian Airlines flight ET302, 10 March 2019. Dated 23 December 2022. The report is formally contested; the United States National Transportation Safety Board issued comments and a press release on 27 December 2022 objecting that its comments were not appended as required, and disputing elements of the technical analysis. https://www.ntsb.gov/news/press-releases/Pages/NR20221227.aspx

49. United States National Transportation Safety Board. *Collision between vehicle controlled by developmental automated driving system and pedestrian, Tempe, Arizona, March 18, 2018*. Highway Accident Report NTSB/HAR-19/03, adopted 19 November 2019. https://www.ntsb.gov/investigations/AccidentReports/Reports/HAR1903.pdf

50. International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC 42001:2023, Artificial intelligence management system*. December 2023.

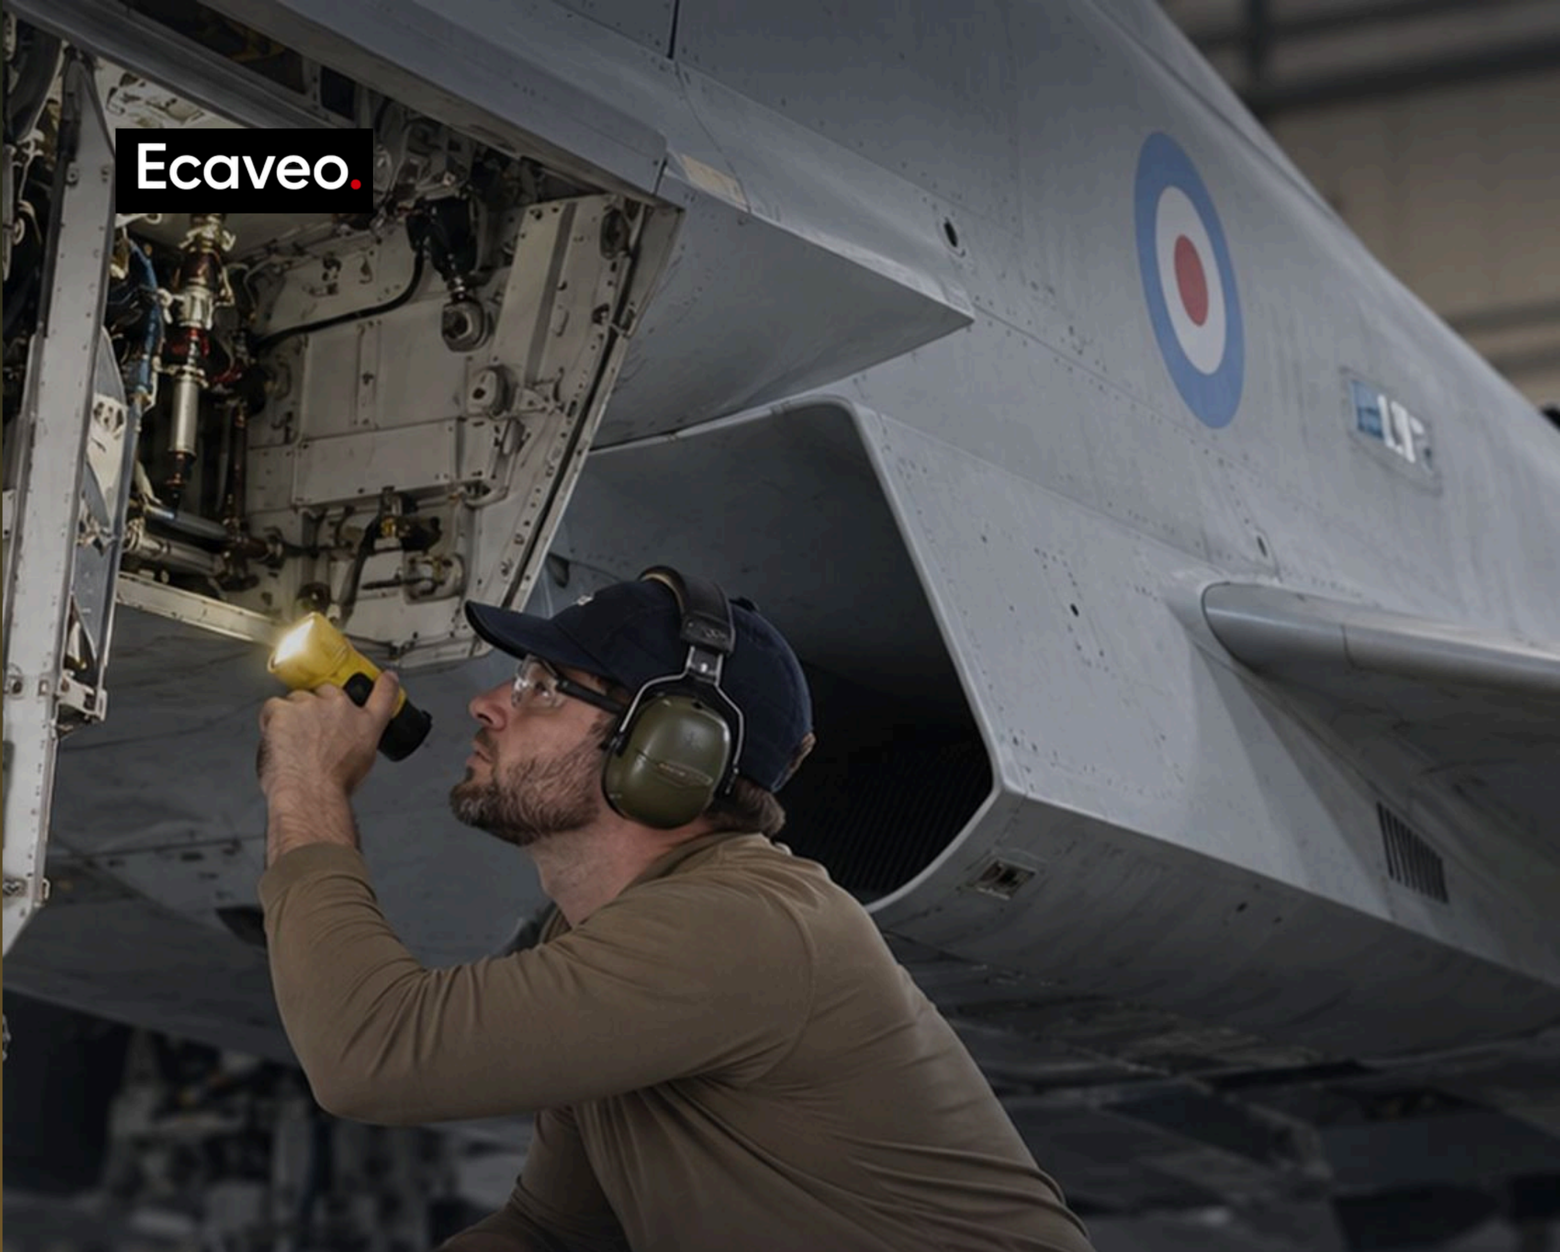

Ecaveo.

ECAVEO RESPONSIBLE AI SERIES

Commensurate Confidence

Ecaveo Responsible AI Series, Aerospace and Defence. Version 1.0, July 2026.

ALSO IN THIS SERIES

The Known Safe State. Public services. Partial Understanding. Financial services.

The Protected Constraint. Operations. Verified Authority. Law firms.

THE AUTHOR

Paul Forrest is a fractional Head of AI working with private equity and venture capital portfolio businesses on the deployment of artificial intelligence.