



Ministry
of Defence

Laying the Groundwork

Responsible AI Senior Officers' Report

October 2025



Laying the Groundwork

Responsible AI Senior Officers Report 2025

Contents

Contents	1
Introduction	2
Strategic Engagement Beyond This Report	2
MOD’s Responsible AI Policy Approach	2
Progress since 2024	3
From Policy to Practice	3
Responsible AI Senior Officers – Explainers of Terms	3
What do RAISOs do?	3
What do we mean by governance, assurance and Responsible AI?	4
Overall maturity and effectiveness of the RAISO framework	5
RAISO Statements Findings	5
Embedding AI Governance and Leadership across Defence Organisations	6
Culture, Workforce and Skills Development	6
Culture	6
Access to Talent and Subject Matter Experts	7
Recruitment and Training	7
Dedicated RAISO Support Roles	7
RAI Risk Management	8
AI Landscape	8
RAI Risk Management	9
Assuring Complex AI Applications	9
Relationship with Industry	9
Summary	10
Glossary	11

Introduction

This report sets out a summary of the work being done across the Ministry of Defence to ensure an ambitious, safe and responsible approach is taken to the development, integration and use of Artificial Intelligence (AI) in the department. It details the groundwork that is being laid to drive organisational change, ensuring the right governance mechanisms, risk management, training and accountabilities are in place to specifically address the benefits and considerations associated with different AI use cases. It constitutes a snapshot in time, providing a high-level overview of some of work to date which the MOD makes public as part of our transparency commitment.

We have a duty to use the best technologies to protect the UK's and our allies' interests, and we also have a duty to use these technologies in line with our legal obligations and the values of the public we serve. The [Strategic Defence Review 2025 \(SDR\)](#) recognises that as part of increasing our warfighting readiness to deter threats and strengthen our security, innovations in AI and digital warfare will be key to making our Armed Forces stronger and safer. As we align delivery against the directives in the SDR, we place emphasis on good governance as it is fundamental to delivering a societally acceptable culture of Responsible AI (RAI).

Our 2022 case for developing and using AI where it delivers clear benefits or more ethical outcomes has only strengthened. In a time of global competition and limited resources, Defence must maximise effectiveness and efficiency through AI. The SDR commits us to world-leading innovation in autonomous systems. Yet, AI also brings risks and concerns about human impact. By proactively addressing challenges that arise from AI integration, we can strengthen public trust, safeguard operational effectiveness, and set a responsible standard for AI-enabled defence.

Strategic Engagement Beyond This Report

This report does not address wider impacts of AI adoption but focusses specifically on assessing the Ministry of Defence's organisational maturity in implementing RAI practices. Our understanding of AI's benefits, risks, and harms, including those affecting civilians, forms a foundational element of the UK's policy approach¹: The department continues to actively engage in international dialogues on RAI, lethal autonomy, and related strategic challenges, including through the UN, to promote shared understanding and to advance norms on the application of International Humanitarian Law to autonomous weapons systems.

MOD's Responsible AI Policy Approach

In 2022 the MOD published five AI ethical principles in the [Ambitious, Safe and Responsible policy paper](#), setting clear organisational intent and ambition for the adoption and exploitation of AI, backed up with defined roles and responsibilities. In 2024 the Ministry of Defence (MOD) published the [Dependable Artificial Intelligence \(AI\) Joint Service Publication 936 \(JSP 936\)](#) to set more detailed policy directives for RAI across Defence.

The MOD established its AI Ethics Advisory Panel in 2021 to provide expert advice, scrutiny and challenge across the full span of principles, policies and frameworks relevant to the delivery of ethical AI outcomes within Defence. It brings together experts (including critical

¹ For more information on this the [Government Response to the House of Lords AI in Weapon systems Committee Report](#) can be accessed.

perspectives) from Defence, academia, industry, and civil society. The panel has now met ten times. It was instrumental in the development of the MOD's AI Ethical Principles, advised on the development of the Defence AI Strategy 2022 and has provided constructive challenge on our developing approach to operationalising AI ethics, including by reviewing draft policy, proposing external best practice, and reviewing real life case studies.

Progress since 2024

In response to the publication of JSP 936, each MOD component organisation (CO; see list on page 4) nominated a Responsible AI Senior Officer (RAISO). RAISOs are senior Defence leaders who ensure their CO adopts and governs AI technologies responsibly, in line with MOD policy. They champion ethical AI use and embed governance within existing structures. RAISOs also engage with a wider community of practice and experts to manage AI risks and support Defence-wide alignment on RAI. As set out in JSP 936, each CO must provide Statements of AI Ethical Assurance which are underpinned by appropriate, auditable evidence.

From Policy to Practice

The publication of JSP 936 has positioned the MOD as a leader in RAI, prompting several COs to update internal policies and enabling delivery teams to progress AI use cases with greater confidence. Its flexible, risk-based approach has been welcomed by both MOD teams and industry, and supports tailored, evidence-led assurance. As AI technologies and our understanding of their adoption evolves, the policy will also be iterated.

AI technologies are becoming ubiquitous and pervasive. They can no longer be treated as experimental and therefore we need to transition their management into that of any other technology that our COs maintain. However, each AI development team and user group needs to contextualise the principles and policy directives to their own unique use case.

Defence has made a confident start in transforming our COs to harness drones, data, and digital warfare to make our Armed Forces stronger and safer. The goal is to gradually embed AI into everyday decision-making and operations – where that makes sense and delivers advantage. As key enablers like data infrastructure and governance continue to evolve, AI teams are actively shaping best practice by putting existing policy into practice and helping build a robust support ecosystem for the future.

Overall, the RAISO statements demonstrate that the level of governance matches the level of AI activity – as our COs increase their AI adoption, existing governance mechanisms are dynamically adapted. They reflect not only policy intent but a shared commitment to RAI oversight. These statements establish a baseline – highlighting progress, remaining challenges, and a collective ambition to make RAI standard practice.

Responsible AI Senior Officers – Explainers of Terms

What do RAISOs do?

The RAISO oversees the responsible adoption and governance of artificial intelligence (RAI) and submits annual statements for organisational assurance. These statements are supported by auditable evidence to make sure COs comply with MOD AI policies and principles.

The RAISO is responsible for making sure that their CO has the right processes, policies and escalation methods in place to demonstrate that they have addressed the considerations

arising under each of the AI ethical principles in respect to AI-enabled systems. RAISOs do this by creating a governance mechanism through delegations to subject matter experts within the relevant area to address the issues raised by capability adoption, overseeing training and creating cultural change.

RAISOs are increasingly supported by a growing network of technical experts, change leaders, project managers and operational colleagues. The evolution of these communities will be essential to sustaining momentum and embedding RAI practices across Defence.

A RAISO has been nominated in each of the following Defence Commands and COs: Army, Royal Navy, Royal Air Force, Cyber and Specialist Operations Command (formerly Strategic Command), Department of State, Defence Equipment and Support, the Defence Science and Technology Laboratory, Defence Business Services, Defence Safety Authority, Defence Infrastructure Organisation, UK Hydrographic Office, Ministry of Defence Police, Sheffield Forgemasters, Single Services Regulations Office, Oil and Pipeline Agency, Submarine Delivery Agency, and AWE Nuclear Security Technologies.

What RAISOs do not do:

While RAISOs play a critical leadership role in embedding RAI across Defence, they are not responsible for technical delivery, system-level assurance, or direct operational deployment of AI capabilities. They do not act as developers, testers, or certifiers of AI systems. RAISOs do not own all organisational levers and can delegate some responsibilities to subject matter experts who oversee AI development and use in their respective areas. As is the case with all other capabilities and services, for AI-enabled systems, most risks are managed at the local level with SROs escalating more significant risks as they arise. RAISOs ensure that appropriate governance structures, escalation pathways, and assurance processes are in place, working in concert with delivery teams, technical experts, and central enablers such as the [Defence AI Centre \(DAIC\)](#) and the Defence AI and Autonomy Unit (DAU – the AI policy and strategy team in the Department of State). The RAISO role is strategic and facilitative, not operational or technical.

What do we mean by governance, assurance and Responsible AI?

Governance describes the frameworks, processes and structures required to make sure that Defence meets our objectives and complies with relevant laws and regulations.

Assurance is about gathering evidence and structuring it into a logical argument. This helps decision makers understand the level of confidence that can be placed in an algorithm or system. Without appropriate assurance, we run the risk of using systems too soon when they are still unsafe, or too late and missing valuable opportunities. The [Dstl Biscuit Book – Assurance of AI and Autonomous Systems](#) provides a more detailed overview on this topic.

Any AI governance and assurance we set needs to build on roles and responsibilities that we already have. As AI development creates different types of risk, we might need more governance in some areas. It's up to RAISOs to make sure this is done in a way that meets the vision for the use of AI in their CO, supported by central advice and other enablers provided by the DAIC and DAU.

Responsible AI (RAI) is the deliberate and accountable use of AI that aligns with legal obligations, ethical principles, and organisational values, ensuring robust governance and

effective risk management. Deploying AI effectively and responsibly – whether in the back office or at the front-line – is a complex process to ‘get right’ because the AI must work within sociotechnical systems where human judgement, organisational culture, and machine capabilities interact. Success depends on effective human-machine teaming and clear governance to support context-appropriate human involvement, not just technical performance. Context appropriate human involvement requires focus on the role of the human as the accountable actor and therefore frames the necessity of human interventions to the servicing to their legal obligations in relation to their particular role in the governance framework, from development, to assurance, to use. For this reason, RAISOs and other leaders in their COs work to ensure the right level of governance and attention is paid to each endeavour.

Embedding our AI policies and processes helps to accelerate our responsible AI adoption because they give clarity and confidence that AI-enabled capabilities have been developed in a robust way. The JSP 936 publication has helped industry engage with MOD as our industry partners are able to understand MOD’s clear AI policy requirements).

Overall maturity and effectiveness of the RAISO framework

“The challenge to successfully acquire, manage, assure, and integrate AI in a responsible, streamlined, and effective manner across a broad spectrum of equipment and business applications should not be understated. Should we be able to do this and deliver effective AI for us and our mission partners, we could rightly claim ourselves as a world leading organisation in AI.” DE&S

MOD’s implementation of RAI is at a formative stage. Momentum is building, with early successes evident in COs that combine strong leadership with technical capability. Foundational structures are being laid, and gaps are being closed, particularly in skills, data management and governance integration. Scaling AI responsibly will require sustained investment, and Defence recognises that embedding RAI as business-as-usual demands cultural change supported by the right enablers.

The RAISO framework and organisational assurance process have played a key role in advancing responsible AI adoption while encouraging innovation. The introduction of RAISO assurance statements has helped COs assess their current baseline, identify what works, and surface persistent barriers. These insights form a critical evidence base for prioritising future investment and shaping strategic direction. They also shape the ongoing development of assurance processes and methodologies which are a joint effort between many COs.

RAISO Statements Findings

Summary of AI Governance in Defence organisations

“The Army understands its use of AI must be carefully planned, but governance should not be overly bureaucratic, so a balance is essential.”

Embedding AI Governance and Leadership across Defence Organisations

Most MOD COs reported that AI features in their organisational strategy to some extent, with several aspiring to become RAI flagship organisations. The leadership role of RAISOs in setting a clear vision and holding their CO to account is critical.

Some COs have used wider organisational change processes to embed the management of RAI principles under existing groups and have specifically invested in integrating AI governance, embedding JSP 936 directives into existing assurance processes. Several COs are using groups around the Chief Technology Officer function to cohere AI, and a few have dedicated AI centres which act as a single point of entry and ensure a cohesive approach to engagement. A minority of COs reported that they are not yet exploiting AI and will adopt the DAIC's guidance on AI Assurance frameworks when AI use cases emerge.

Some COs tailor the centrally developed AI Assurance Framework to their specific operational contexts. Doing so enables more relevant, streamlined assurance processes that align with active use cases, making it easier for delivery teams to embed RAI principles by design rather than retrofitting compliance.

RAISO statements indicated that the transfer of AI-enabled capabilities – and therefore responsibility – from R&D and acquisition teams to operational COs will increase. Building on early examples, relevant COs are formalising coordination groups to ensure continuity of oversight, shared visibility of programmes, and a consistent approach to managing RAI risks across the entire capability lifecycle.

Several COs play a pivotal role in enabling wider AI adoption through, for example, setting AI policy and steering commercial and R&D AI activities – with central coherence enabled by DAU/DAIC. While some of these bodies explicitly recognise their enabling role and are actively investing in the growth of their AI capabilities, others are still in the early stages of resourcing and formalising their AI functions. These AI-enabling COs continue to work with central functions like the DAIC and DAU to clarify what stakeholders need (for example, tailored advice on AI's regulatory impact) and scale their expertise accordingly.

Culture, Workforce and Skills Development

Culture

Most COs reported that staff are keen to explore the benefits of AI. Often, there are pockets of accelerated AI development within a CO, whilst other areas remain under-exploited. A shared challenge arises in how COs identify, resource and manage the increasing demand for AI, balancing the experimentation of technology with the delivery of robust products that move into live service. The RAISO statements also recognise that complexity increases when embedding new AI-enabled tools within existing platforms or importing R&D from other COs. Understanding and adoption of RAI practices is growing, and RAISO statements highlighted ongoing efforts to accelerate progress and address remaining gaps.

Access to Talent and Subject Matter Experts

Recruitment and Training

Beyond establishing sustainable, direct support to RAISOs, MOD faces the tri-fold task of a) recruiting external AI experts, b) training some existing personnel to a high level of expertise, and c) upskilling the rest of the workforce to become sufficiently AI literate.

To become an intelligent customer of AI, Defence must strengthen its ability to discern meaningful engagement with responsible practices from bad practices. Over half of COs have introduced AI ethics guidance and training, and several are actively disseminating tailored materials. The DAIC continues to support this effort by developing central guidance and assurance frameworks, drawing on cross-government and allied expertise. For more details on this, see [Section 5](#) (pages 22-24) in the Government's response to the House of Commons Report on Developing AI Capacity and Expertise in UK Defence.

Dedicated RAISO Support Roles

RAISOs have drawn on cross-functional expertise – including cyber, legal, ethics, and science and technology – to implement Responsible AI across Defence. In some cases, technical experts support RAISO functions alongside their core roles, which has worked well for select projects. As AI adoption grows, there is recognition that more structured and scalable support will be needed, particularly to strengthen domain-specific AI expertise.

While this ad hoc support has been instrumental in early progress, RAISO statements underscored the value of having dedicated teams in place which support the RAISO directly. This is particularly relevant for larger COs with a growing portfolio of ambitious AI uses cases.

Given the varying scale and complexity of AI portfolios across Defence, a tailored approach is essential. The DAIC and DAU continue to collaborate with RAISOs to develop and share best practice models that reflect organisational diversity while promoting consistent standards.

Early examples have shown that the establishment of Responsible AI Officers (at working level) with clear escalation routes to the RAISO level are a model for success: Akin to establishing safety and mishap programs, RAI officers in military units and COs can serve as local conduits for new information and guidance, promoting broad-based AI literacy, reporting AI incidents, and mitigating AI risks. Through docking into centrally organised Communities of Interest (i.e. through the DAIC), RAISOs ensure that a group of dedicated points of contact in their COs is empowered to cohere the dissemination of new directives, training, guidance. This approach has also been recommended in the report on [AI for Military Decision Making](#), by the Center for Security and Emerging Technology.

Multidisciplinary and diverse AI development teams

RAISO statements reflected that delivering AI safely, ethically and effectively requires more than technical expertise. Successful use cases consistently demonstrate the importance of multidisciplinary teams, including cyber, data legal, ethical, commercial and AI technical specialists, working alongside project managers and RAISOs to ensure coherence and responsible implementation.

The Strategic Defence Review highlights the need for a coordinated effort to build and sustain the skills required for effective and responsible AI deployment. There was strong recognition in all RAISO statements that access to the right talent and expertise will be crucial to achieving

the MOD's ambitions for AI adoption. Defence is actively building a strong pipeline of AI talent by establishing Suitably Qualified and Experienced Personnel (SQEP) teams, including blended Crown and contractor groups, and is accelerating coordinated workforce and training plans to ensure we attract, develop, and retain the expertise needed for responsible AI deployment.

Senior Leadership

RAISO statements were clear that Defence needs to avoid both underutilisation of promising capabilities and exposure to low-quality AI offerings. Specific responsibility was placed on leaders to build sufficient understanding of AI's risks and benefits, to set the right expectations and put in place appropriate support to enable successful delivery. A strong emphasis on the importance of continued senior leadership training in AI and RAI practices characterised most statements.

Wider Workforce Upskilling

Whilst specific attention is being paid to expert AI development teams, most statements noted that the wider organisations' workforces need to reach a basic, sufficient competency level in AI to a) use any AI-enabled tools responsibly, and b) become a more intelligent customer by understanding where AI can be effectively deployed. To support this, COs leverage central AI training and guidance, noting that for some areas, RAISOs need to consider more specific and tailored upskilling needs, depending on the nature of AI use cases. Central programmes include:

- MOD's Digital Skills for Defence (DS4D) programme, which is delivering a critical capability uplift across Defence, equipping leaders, specialists, and the wider workforce with the skills needed to deploy AI and digital technologies responsibly.
- The [AI Ethics Playing Cards](#) can be accessed and serve as a useful tool for prompting AI development team discussions about applying AI ethics in practice in a Defence context.

In summary, while there is a growing cadre of AI professionals in some areas, the most significant challenge identified across all RAISO statements is increasing the availability of SQEP to support Defence's commitment to RAI adoption. Addressing this will be critical to embedding AI as a sustainable and strategic capability.

RAI Risk Management

AI Landscape

COs employ AI across a vast spectrum of use cases, spanning both operational and business domains. Most COs have successfully developed a strong understanding of their AI landscapes. They not only articulate a clear vision of the benefits they aim to realise but also demonstrate proactive and holistic awareness of the risks they need to mitigate, recognising the legal, safety, and ethical dimensions. Related activity includes COs commissioning research to deepen understanding of RAI maturity and updating internal policies to incorporate AI considerations.

As COs increase their use of AI, RAISOs are building on their mechanisms for tracking their area's AI landscape in a pragmatic and futureproofed way, using or uplifting existing internal governance tools.

RAI Risk Management

The RAISO statements collectively demonstrate a strong and nuanced understanding of the key benefits and risks associated with Defence AI adoption, particularly across legal, safety, and ethical dimensions.

Overall, in terms of risk management, JSP 936 policy directs that COs should use and adapt their *existing* risk management processes to address AI-specific considerations.

To ensure a holistic assessment of AI benefits and risks is made, MOD developed a defence AI ethics risk assessment toolkit which gives guidance on how ethical considerations might impact, increase or alleviate traditional risks (such as health and safety, financial and reputational) but could also introduces new ones which have to be balanced against the benefits being pursued.

Local implementation of AI ethics management varies across COs but is appropriate to ethical risks held. To maintain the right balance between speed of development and maintaining safeguards, risks are managed at the lowest appropriate level. However, UK MOD policy is clear that for AI applications that are inherently novel, higher risk or contentious, consideration or sign-off at very senior levels (including ministerial) may be required.

Defence continues to embed AI considerations into established risk management frameworks, reflecting its growing maturity in the domain. As AI capabilities become increasingly integral to Defence operations, they will be treated as core assets rather than exceptional cases. Existing escalation pathways for ethical, legal, and safety concerns will be reinforced to ensure that RAI risks are identified and addressed early. This will empower personnel to raise concerns with confidence and support timely, informed decision-making across the organisation.

Assuring Complex AI Applications

As Defence continues to explore the potential of AI across a growing number of operational domains, including high-risk and complex use cases, the importance of robust assurance mechanisms is recognised in RAISO statements, highlighting that Defence is committed to ensuring that ethical, legal and safety obligations are upheld across all applications.

While progress has been made, translating high-level policy into actionable guidance for specific AI use cases and operational contexts remains an ongoing task. Upskilling at senior and official levels is essential to bridge this gap and enable informed decision-making. Defence acknowledges the need for timely, expert advice to support RAISOs, Senior Risk Owners and development teams working on sensitive projects.

Relationship with Industry

Most COs have sought out external challenge and advice during AI development. There is recognition that Defence will need to continue to seek external challenge to ensure we remain alert to any blind spots and avoid becoming complacent in our approach to RAI.

RAISO statements highlighted several constructive supplier relationships, and the RAISO and MOD AI communities have benefitted from engagement with innovative and forward-leaning industry partners. Some suppliers are taking an exceptionally conscientious approach to RAI implementation, embedding ethical, legal, and safety considerations from the outset. The MOD seeks to continue to learn from these thought leaders to raise the bar for AI assurance.

As Defence accelerates its adoption of AI, it is committed to working with industry to ensure that all solutions meet the highest standards of legal, ethical and operational integrity. Recent organisational feedback highlights a growing need for suppliers to demonstrate clear compliance with legal and ethical frameworks and to engage with assurance processes that reflect the complexity and sensitivity of Defence use cases.

To support this, COs can apply tools such as DSIT's supplier questionnaires and model cards into procurement processes, aligning them with Secure by Design principles. These steps are part of a broader ambition to lead by example - setting clear expectations, building trust, and ensuring that AI adoption is both responsible and operationally effective.

However, challenges remain. Infrastructure limitations, unclear development pathways, and variability in supplier capability can complicate delivery. Additionally, COs have noted a pressing need for more consistent contract terms and access to SQEP to support robust assurance. Without these, there is a risk of fragmented efforts and duplicated work, particularly where individual COs lack the maturity to act as intelligent customers. This not only impacts quality and efficiency but can also lead to supplier frustration and lack of user adoption of AI tools.

To address this, Defence is considering how to better coordinate supplier engagement across the enterprise. DAIC and DAU will work with the RAISO community to enable forums for sharing learning and supplier experiences. The central mechanism for this is the AI Expert Group (a community of AI industry and academic stakeholders) which has already been leveraged to share lessons with RAISOs on external AI governance examples. This will help reduce duplication, improve consistency, and ensure that Defence COs are equipped to demand high-quality, assured AI solutions.

RAISOs are encouraged to work closely with commercial teams to develop contract terms that reflect RAI expectations and to use existing supplier relationships to continue to drive up standards. Defence remains committed to fostering a collaborative environment where industry partners are supported in delivering innovative, trustworthy AI capabilities that meet Defence's operational needs.

Summary

In just under a year, the RAISO community has laid a powerful foundation for RAI across Defence, transforming policy into practice and ambition into action. By acting on the report's recommendations, the MOD can lead by example - setting the standard for ethical, assured, and operationally effective AI. DAU and DAIC will work with RAISOs over the coming months to prioritise and implement these recommendations.

Glossary

CO – Ministry of Defence Component Organisation

DAIC – Defence AI Centre

DAU – Defence AI and Autonomy Unit

RAI – Responsible AI

RAISO – Responsible AI Senior Officer

SDR – Strategic Defence Review

SQEP – Suitably qualified and experienced personnel

The report has been written by humans. AI tools have been used for spell checks and summarisation which have been reviewed by humans.