



GC REAIM
GLOBAL COMMISSION ON RESPONSIBLE
AI IN THE MILITARY DOMAIN



تريندز للبحوث والاستشارات
TRENDS RESEARCH & ADVISORY




Foreign, Commonwealth
& Development Office

GC REAIM Expert Policy Note Series

A systems engineering lifecycle approach to responsible AI

Zena Assaad and Adam Hepworth

May 2025



POWERED BY



The Hague Centre
for Strategic Studies

GC REAIM Expert Policy Note Series

A systems engineering lifecycle approach to responsible AI

Authors: Zena Assaad and Adam Hepworth

May 2025

Cover photo: [unsplash](#)

The Global Commission on Responsible Artificial Intelligence in the Military Domain (GC REAIM) is an initiative of the Government of the Netherlands that was launched during the 2023 REAIM Summit on Responsible Artificial Intelligence in the Military Domain in The Hague. Upon request of the Dutch Ministry of Foreign Affairs, the Hague Centre for Strategic Studies acts as the Secretariat of the Commission.

The GC REAIM Expert Policy Note Series was funded by the Foreign, Commonwealth and Development Office (FCDO) of the United Kingdom. GC REAIM Experts maintained full discretion over the topics covered by the Policy Notes. The contents of the GC REAIM Expert Policy Note series do not represent the views of the Global Commission as a whole. The Policy Notes are intended to highlight some key issues around the governance of AI in the military domain and provide policy recommendations. This Policy Note was released during the Commission Meeting in Abu Dhabi, hosted and graciously supported by Trends Research and Advisory.

© The Hague Centre for Strategic Studies. All rights reserved. No part of this report may be reproduced and/ or published in any form by print, photo print, microfilm or any other means without prior written permission from HCSS. All images are subject to the licenses of their respective owners

HCSS
Lange Voorhout 1
2514 EA The Hague

Follow us on social media:
@hcssnl

The Hague Centre for Strategic Studies
Email: info@hcss.nl
Website: www.hcss.nl



تريندز للبحوث والاستشارات
TRENDS RESEARCH & ADVISORY



**The Hague Centre
for Strategic Studies**

1. Introduction

The opacity of AI-enabled systems employment in the military domain is encouraging ever more nuanced discussion around the need for new and innovative approaches to safely and responsibly develop and implement these technologies. These discussions can often overlook existing systems engineering processes for complex system development, which are transferable to AI-enabled systems. Systems engineering is the process of designing, integrating and managing a system over the course of its lifecycle. Here, a system is taken to mean a composition of multiple parts.¹ Wholeness is a characteristic of systems, because properties and behaviours are derived from the interactions between components, they do not emerge in isolation.² These are referred to as emergent properties, because they manifest at the systemic level.

Complex systems do not operate in isolation of the broader ecosystem they sit within. Take for example an aircraft. The system is not confined just to the physical aircraft. Rather, it includes the designers, manufacturers, regulators, maintainers, pilots, etc. This is because an aircraft cannot function or operate in the absence of these broader components which are spread across its lifecycle. AI-enabled systems are complex systems. They are designed, developed and operated through an interaction of processes, procedures and decision making across the entirety of their lifecycle. As such, the system should be viewed holistically, capturing that broader lifecycle.

The UN's high level advisory body on AI released a final report, *Governing AI for Humanity*, which makes reference to "...promoting responsible life cycle management of AI applications in the security and military domain."³ The report also advocates for developing responsible AI guidelines, for developers and designers of AI technologies, as well as for their users and all actors involved throughout their lifecycle.

This brief advocates for the use of a foundational systems engineering lifecycle model as a contribution to responsibly developing and implementing AI-enabled systems in the military domain. While this paper focuses on technical and procedural accountability, we acknowledge broader ethical and legal considerations as essential elements in achieving a comprehensive responsible military AI outcome.

A high-level use case is presented and applied to the proposed model to demonstrate how it would apply in practice. The significance of the use case in this brief is twofold. Firstly, it demonstrates the effectiveness of the proposed lifecycle model. Secondly, it highlights the importance of context specificity. The use of AI capabilities spans a diversity of contexts, domains and applications. In order to maintain proportionality, context specificity is required in approaches to responsible AI.

¹ Andrew P. Sage and Christopher D. Cuppan, 'On the Systems Engineering and Management of Systems of Systems and Federations of Systems', *Information Knowledge Systems Management* 2, no. 4 (1 November 2001): 325–45, <https://doi.org/10.3233/IKS-2001-00045>.

² Gail G. Whitchurch and Larry L. Constantine, 'Systems Theory', in *Sourcebook of Family Theories and Methods: A Contextual Approach*, ed. Pauline Boss et al. (Boston, MA: Springer US, 1993), 325–55, https://doi.org/10.1007/978-0-387-85764-0_14.

³ United Nations, *Governing AI for Humanity* (United Nations, 2024), <https://doi.org/10.18356/9789211067873>.

2. Lifecycle model for responsible AI

There exists a number of responsible AI initiatives and frameworks, with a majority focusing on civil applications of AI. These vary from independent organisations developing internal principles, academic researchers proposing methodologies or frameworks and nonprofit organisations proposing principles for the responsible development and use of AI. Examples of such frameworks and initiatives can be found in Gollner et al (2023).⁴

Despite the growing body of knowledge in this field, convergence on a definition of responsible AI remains elusive. The terms *responsible*, *responsibility* and *ethical* are commonly used interchangeably among the literature. Additionally, most ethical initiatives make note of *responsible* and *responsibility*, and vice versa. These terms have become synonymous with one another, making efforts to delineate them difficult.

For the purpose of contextual clarity within this brief, the terms *responsible* and *responsibility* are considered synonymous and are taken to mean *attributions of responsibility, in terms of decisions and actions, and accountability for those decisions and actions*.⁵ For our purposes here, responsible AI refers to *AI-enabled systems developed and implemented such that actions are traceable and accountable across the system's lifecycle*. This brief adopts a human centric understanding of responsible AI, with the decisions and actions being those of humans across a system's lifecycle.

This brief focuses on technical aspects of accountability, with *ethical* concepts considered outside the scope of this work. While it is commonly associated with the notion of responsible and responsibility, the concept of ethics is considered too broad and subjective to define and capture in this context.

While there are many responsible AI principles and approaches among the literature, few have a military focus.⁶ A recent IEEE report developed by the IEEE SA research group on issues of autonomy and AI in defence systems presented a framework for human decision making through the lifecycle of autonomous and intelligent systems in defence applications.⁷ The report adopts the approach of examining the holistic lifecycle of a system and identifying who should be responsible and accountable at every stage across that lifecycle. The report focuses on the accountability of potential problems, conflicting guidance or legal jeopardy.

⁴ Sabrina Goellner, Marina Tropmann-Frick, and Bostjan Brumen, 'Responsible Artificial Intelligence: A Structured Literature Review', 2024, <https://doi.org/10.48550/ARXIV.2403.06910>.

⁵ Zena Assaad and Christine Boshuijzen-van Burken, 'Ethics and Safety of Human-Machine Teaming', in *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, TAS '23 (New York, NY, USA: Association for Computing Machinery, 2023), 1–8, <https://doi.org/10.1145/3597512.3600205>.

⁶ Jan Maarten Schraagen, 'Responsible Use of AI in Military Systems: Prospects and Challenges', *Ergonomics*, 2 November 2023, <https://www.tandfonline.com/doi/full/10.1080/00140139.2023.2278394>.

⁷ Sten Allik et al., 'A Framework for Human Decision-Making through the Lifecycle of Autonomous and Intelligent Systems in Defense Applications', *A Framework for Human Decision-Making through the Lifecycle of Autonomous and Intelligent Systems in Defense Applications*, October 2024, 1–63.

While not explicitly addressing *responsibility*, the contents of the report embody measures towards attributions of accountability, which define the notion of responsibility as presented in this brief. The framework within this report is used as the basis of the responsible AI framework of this brief, with a few amendments. The first is a simplification of the lifecycle from nine to four stages. The approach presented here is further simplified to reflect the technical focus of this work, in contrast to the legal focus of the IEEE framework. This brief builds on the existing corpus of knowledge, including the influence of core standards for the AI lifecycle⁸ and ethical design standards.⁹

The four lifecycle stages adopted within this brief are:

1. **Plan:** The planning phase includes defining the scope and determining the objectives of a system.
2. **Design:** The design phase involves outlining specific requirements, architectures, functions, interfaces, etc. of the system.
3. **Develop:** The development stage involves the actual building of the system, which includes test and evaluation¹⁰ and verification¹¹ and validation¹².
4. **Deploy:** The deployment stage involves integrating the system into operations, which includes maintenance throughout the life of the system.

The second amendment made is adopting only three of the five ongoing activities. The IEEE framework presents five activities which need to occur repeatedly at each lifecycle stage. In this brief, we have adopted the following three:

Activity 1: Evaluation of legal, regulatory and policy requirements.

Activity 2: Considering the human: training, education, and human-system integration.

Activity 3: Risk assessments.

Note, the first activity has been amended from its original wording, replacing *ethical* with *regulatory* and *concerns* with *requirements*. Ethical measures are considered out of scope of this brief and the term requirements has been used to reflect the technical focus of this approach. The proposed simplified human centric lifecycle model for responsible AI is presented in figure 1.

⁸ 'IEEE Standards Association', IEEE Standards Association, accessed 6 May 2025, <https://standards.ieee.org/ieee/7000/6781/>.

⁹ 'ISO/IEC 5338:2023', ISO, 2023, <https://www.iso.org/standard/81118.html>.

¹⁰ Test and evaluation involves assessing the performance, reliability, and safety of a system. It involves systematically examining and validating these items to ensure they meet specified requirements and perform as intended.

¹¹ Verification refers to the evaluation of whether or not a product, service, or system complies with a regulation, requirement, specification, or imposed condition. It is often an internal process.

¹² Validation refers to the assurance that a product, service, or system meets the needs of the customer and other identified stakeholders. It often involves acceptance and suitability with external stakeholders.

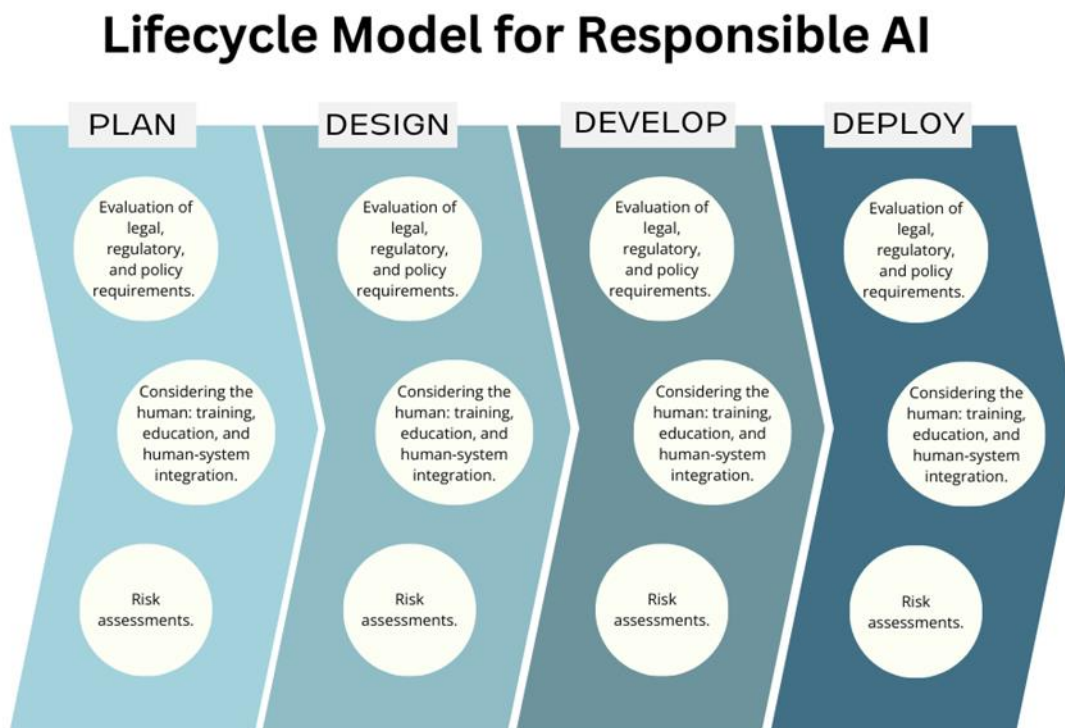


Figure 1. Lifecycle model for responsible AI

3. Use Case: AI-enabled command and control

AI-enabled command and control (C2) offers a transformative approach to warfare by leveraging autonomous and collaborative decision-support systems to enhance and augment human centric military operations. In this use-case, the system relies on integrating a network of heterogeneous teaming systems across multiple domains, generating a force-multiplying effect that accelerates human decision making. This idea relies on the notion of battle management systems, with a collection of AI-enabled systems identifying possible actions and responses, ensuring commanders have real-time access to high-fidelity information, while retaining human centric command over decisions and risk appetites within the operational environment.

This vignette focuses on a tactical organisation supported with an AI-enabled C2 capability. The commander may have a range of autonomous capabilities including uncrewed aerial systems, uncrewed ground vehicles, and crewed platforms, all interconnected through the integrated battle management system. The AI-enabled C2 system supports the commander in finding and identifying potential threats, tracking movements within the battlespace and augmenting human-centric decision-making processes. Prior to the mission, the commander conducts a detailed analysis of the battlespace. Aided by an embedded decision-support system, the commander explores different courses of action to develop robotic and autonomous systems mission profiles. Systems of control are continuously refined to ensure intent alignment and to meet operational and legal obligations. Thorough assessments provide a detailed view of outcomes from different actions.

After the commencement of a mission, areas of interest and priority intelligence requirements are assigned to human-machine teams. Identification of possible threats are cross-confirmed with autonomous and human team-mates, prior to a comprehensive human centric assessment being provided to the commander for consideration. Importantly, if no decision is made by the commander regarding a possible action, no further action is taken within the battle management system.

Upon a threat being identified and cross-confirmed, it is subsequently verified and validated by the commander, including aspects of human led legal review. The significance of the threat is evaluated against the performance and monitoring systems developed during planning phases. The commander decides if the threat is of high priority, with the AI-enabled C2 system proposing specific actions at the time of decision. Upon authorisation of a specific course of action, continuous real-time oversight and monitoring of both the vehicle, the immediate area and the effect method are undertaken. The commander ensures that the threat remains valid, poised to adapt actions as necessary. Post-action, the AI-enabled C2 system tasks an autonomous reconnaissance team to confirm the status of the threat, validating estimates through a battle assessment and updating system records, initiating further course of action analysis and development. Continuous performance monitoring activities assure commanders of expected versus actual outcomes, enhancing system confidence.

4. Systems Engineering Approach to Use-Case Based Responsible AI

We now apply the lifecycle model for responsible AI to the use-case presented in the previous section to demonstrate how the model can be used in practice. Each discrete activity requires clearly determining accountabilities and responsibilities, as well as ensuring audit artefacts such as records, logs and key decisions across all phases are maintained to promote system transparency.

4.1 Lifecycle phase 1: Plan

The planning phase includes defining the scope and determining the objectives, constraints and limitations of use for a system. For the AI-enabled C2 system, this phase may entail determining the boundaries of what the C2 system can and cannot do. It would also involve establishing what the commander's role would be in this system, as well as defining the relationship between the commander (human) and AI-enabled C2 system (machine). Required interfaces and safety testing may also be determined during this phase as anticipatory planning of requirements is commonly captured during the early stages of a lifecycle.

Applying the three ongoing activities to this phase for the C2 system will involve mostly anticipatory planning and forecasting of what will be needed. This is highlighted in the table below.

Activity	Outcome
Evaluation of legal, regulatory, and policy requirements	Identify what existing laws, regulations and policies will apply to the C2 system. These should encompass requirements for the design and implementation of the system. If measures do not yet exist, consider utilising standards to fulfil this activity. While not legally binding, standards can act as a useful substitute in the interim to help guide responsible development.
Considering the human: training, education, and human-system integration	Determine what training and education will need to be provided for all of the humans involved in this system. This extends beyond the commander and system operator to include maintainers, safety personnel, designers and developers.
Risk assessments	Develop a preliminary risk assessment forecasting the potential risks as well as implementation. This risk assessment should capture timeframes for legal and regulatory compliance, methods of assurance for training and education and other general capability and operational risks. This is in addition to capturing technical risks of the system.

Table 1. Ongoing activities applied to the planning lifecycle stage.

4.2 Lifecycle phase 2: Design

The design phase involves outlining specific requirements, architectures, functions, interfaces and standards for the system. This could include determining such attributes across multiple platforms and use-cases. Because the C2 system, in this case, integrates autonomous platforms across multiple domains, the design phase will be segmented, with each domain requiring separate considerations. While the planning phase takes a more holistic look at the overall system, the design phase requires a more detailed approach.

Applying the three ongoing activities to this phase for the C2 system will involve discrete work on each specific platform, as well as planning for how these platforms can be safely and responsibly integrated. This includes integration with the human actors in the C2 system, including the commander and other operators. This is highlighted in the table below.

Activity	Outcome
Evaluation of legal, regulatory, and policy requirements	Determine what domain specific laws, regulations and policies apply to each of the platforms. Identify if or how these requirements may impact system integration. Some laws, regulations and policies in one domain may conflict or create flow on requirements between platforms.
Considering the human: training, education, and human-system integration	Identify what education and training is required for each platform to ensure responsible integration. Identify what broader education and training may be needed for human-system integration at the broader C2 system level.
Risk assessments	Update the risk assessment from the planning stage to include the more granular details for each platform. Ensure technical risks are captured.

Table 2. Ongoing activities applied to the design lifecycle stage.

4.3 Lifecycle phase 3: Develop

The development stage focuses on system development, which includes test and evaluation and verification and validation (TEVV) to ensure robust oversight across this phase and the broader lifecycle.¹³ This stage requires these processes to be completed at both the platform and system level. A platform may work at an individual level but fail at a system level given the layered and often modular approach to software development. These issues increase with the number of systems integrated. For the C2 system, this stage may also involve TEVV within the human domain, including human-system interaction and integration. This can be achieved through iterative trial events through a mixture of simulated and real-world events.

¹³ David Helmer et al., 'Human-Centred Test and Evaluation of Military AI', 2024, <https://doi.org/10.48550/ARXIV.2412.01978>.

Applying the three ongoing activities to this phase for the C2 system will involve analysing both the platform and system levels, in addition to how they integrate with one another. It will require a phased approach, starting with the technical development and moving into trial operations of the system to determine its effectiveness in operation. There is also a strong focus on human operators in this phase as they form a critical part of the overall system.

Activity	Outcome
Evaluation of legal, regulatory, and policy requirements	Demonstrate compliance with relevant laws, regulations and policies through TEVV procedures. For the C2 system, these will need to be demonstrated and assured at both the platform and systems level.
Considering the human: training, education, and human-system integration	TEVV should capture the commander and other human operators as they form a critical part of the overall system. The effectiveness of their education and training should be demonstrated and assured through these processes.
Risk assessments	Update the risk assessment from the planning stage to include any additional concerns or anticipated risks. The update should also include reassessing previously captured risks to ensure their likelihood and severity are accurately reflected given the progress of the project and the assurance, or lack thereof, achieved during this phase.

Table 3. Ongoing activities applied to the development lifecycle stage.

4.4 Lifecycle phase 4: Deploy

The deployment stage involves integrating the system for operational use. This stage can be viewed in two parts, operation and maintenance. When transitioning to operational use, requirements around education and training, operational and sociotechnical procedures and processes must be maintained and continuously reviewed. Ongoing maintenance of both software and physical platforms will require through life support and may require an iterative upgrade approach. The physical platform may require scheduled maintenance in compliance with existing regulations and standards. Software requirements may be less explicit and maintenance difficult to forecast. Updates may include, for example, datasets, functionalities, interfaces and model re-certification to ensure data drift is captured.

Applying the three ongoing activities to this phase for the C2 system will involve managing the operation of the system and its maintenance concurrently. The operation and maintenance of the system will each require specific processes and procedures, potentially managed by different teams.

Activity	Outcome
Evaluation of legal, regulatory, and policy requirements	Determine what through life support is needed for each physical platform and who is responsible for managing these requirements. Develop a process for determining and implementing interrelated system maintenance and oversight, such as software systems and integrated payloads. Determine what requirements are needed for the operation of the C2 system in different operating environments.
Considering the human: training, education, and human-system integration	Ensure the people responsible for maintaining the platform are sufficiently qualified and experienced. Ensure there are effective education and training requirements for humans within the C2 system, at both the individual and collective organisation levels.
Risk assessments	This stage will require bespoke risk assessments that capture the operation of the system in specific environments. These risk assessments will focus specifically on operational risks, not project risks.

Table 4. Ongoing activities applied to the deployment lifecycle stage.

For the ongoing activities for each of the lifecycle phases, a person or persons need to be assigned as responsible. Responsible AI necessitates AI-enabled systems be developed and implemented such that actions are traceable and accountable across the system's lifecycle. The lifecycle approach presented in this brief requires the allocation of responsibility for each task to be determined and documented. This is to ensure there is transparency on decision making across the system lifecycle.

5. Conclusion

AI-enabled systems are complex systems which operate through an amalgam of interconnected components. In the case of the C2 case study presented in this brief, those interconnected components span multiple platforms and capabilities, scaling the complexity of these systems. The processes of planning, designing, developing and deploying a complex system cannot be viewed in isolation of one another. The lifecycle of such a system will have flow on effects to how the system operates and performs. The lifecycle model presented in this brief demonstrates a means of achieving responsible AI through traceable and accountable actions and decision making across a system's lifecycle. The human centric approach utilises foundational systems engineering processes, highlighting the effectiveness of such principles to minimise the need for novel process and procedure development for these complex systems.


About the Authors

Dr Zena Assaad

Dr Zena Assaad is a senior lecturer in the School of Engineering at the Australian National University and is also a fellow with the Australian Army Research Centre. She has previously held a fellowship with Trusted Autonomous Systems. Her research explores the safety of human-machine teaming and the assurance and certification of trusted autonomous and AI systems. Dr Assaad is the founder and chair of the Australian national community of practice for UAS and AAM research.

Lt. Col. Dr Adam Hepworth

Lieutenant Colonel Adam Hepworth, PhD, serves as Director of the Robotic and Autonomous Systems Implementation and Coordination Office for the Australian Army. In this role, he leads the advancement of emerging technology, including robotics, autonomous systems, artificial intelligence and autonomy for the Australian Army.



HCSS
Lange Voorhout 1
2514 EA The Hague

Follow us on social media:
[@hcssnl](#)

The Hague Centre for Strategic Studies