

Responsible AI and Autonomous Systems Policy

Lifecycle controls for AI-enabled and autonomous systems based on human responsibility, understanding, bias and harm mitigation, reliability, robustness and security.

DOCUMENT ID	RCDEN-AI-POL-001
OWNER	Managing Director
VERSION	1.0
APPROVED	24 July 2026
NEXT REVIEW	24 July 2027 or sooner following material change
CLASSIFICATION	PUBLIC

APPROVAL Authorised for issue by Anthony Deacon, Managing Director, on 24 July 2026. This controlled copy is reviewed at least annually.

1. Policy statement

RCDEN develops and integrates artificial intelligence and autonomy only for lawful, defined and demonstrably beneficial purposes. Human responsibility and accountability remain explicit throughout design, training, testing, deployment, operation, change and decommissioning.

2. Scope

This policy covers machine learning, computer vision, perception, planning, decision support, autonomous navigation, adaptive control, generative or agentic tools, data models and systems in which AI forms a material component. It applies to internal use and customer deliverables.

3. Responsible AI principles

- Human-centricity: assess intended and unintended effects on operators, affected people, society and the environment across the lifecycle.
- Responsibility: define human authority, control, accountability, risk ownership and escalation before use.
- Understanding: make purpose, operating limits, performance, uncertainty, failure modes and required competence understandable to relevant users.
- Bias and harm mitigation: examine data, models, interfaces and operational context for unfair bias, misuse and foreseeable harm, then apply safeguards.
- Reliability: demonstrate that the system is sufficiently reliable, robust, secure and privacy-respecting for its intended context.

4. Lifecycle governance

1. Define the intended purpose, user, environment, benefit, prohibited uses and accountable owner.
2. Identify legal, ethical, safety, security, privacy, human-factors and operational risks.
3. Control data provenance, lawful use, quality, representativeness, access, labelling and retention.
4. Specify measurable performance, confidence, latency, robustness and fail-safe or degraded-mode requirements.
5. Verify, validate and red-team proportionately using representative and edge-case conditions.
6. Document model, software, dataset, configuration, interfaces, limitations, approvals and evidence.
7. Train users and provide clear operating limits, monitoring, override, recovery and incident procedures.

8. Monitor performance and context drift, control changes and reassess before reuse in a new setting.
9. Withdraw, isolate or revert systems that perform outside approved limits or create unacceptable risk.

5. Human control and autonomous weapons

RCDEN will not design or supply a system intended to select and engage human targets without context-appropriate human involvement and accountable human command. No degree of automation removes human responsibility. Any safety-critical or legally sensitive function requires clearly defined authority, constraints, abort or override arrangements and an evidence-based acceptance decision by the competent duty holder or customer authority.

6. Use of generative AI

Customer, personal, security-classified, export-controlled, proprietary or contract-protected information must not be entered into an AI service unless that service and use are explicitly authorised. AI-generated code, analysis, design and text are treated as untrusted until reviewed, tested and approved by a competent person. Intellectual-property, licensing, hallucination and supply-chain risks are assessed before reuse.

NO AUTONOMY WITHOUT EVIDENCE A successful demonstration is not sufficient assurance. Deployment requires evidence that the approved configuration performs within defined limits in the intended environment, with understood residual risk and accountable human control.

7. Responsibilities and review

The Managing Director owns this policy and reviews it at least annually, and sooner after a material change in law, ownership, operations, workforce, supply chain, customer requirement or risk. Actions arising from reviews are assigned, dated and retained as controlled evidence.

Reference framework

The following sources inform this policy. Contract-specific legal, regulatory, security and customer requirements always take precedence.

- [MOD - Ambitious, Safe, Responsible: AI in Defence](#)
- [MOD JSP 936 Part 1 - Dependable AI](#)
- [Government AI Playbook](#)

Document control

VERSION	DATE	APPROVED BY	CHANGE
1.0	24 July 2026	Anthony Deacon, Managing Director	First controlled issue

THE RC DEN LTD, registered in England and Wales under company number 14639255. Registered office: 20-22 Wenlock Road, London, England, N1 7GU. Policy contact: anthony@rcden.co.uk.