



Responsible AI

Service Definition Document

G-Cloud 15

Status: Definitive

Security Commercial – In Confidence

Version no.1/0

THINKING
BEYOND
LIMITATIONS

We're Kainos

Formed in 1986, we set about finding the best people to tackle the most difficult digital work. We fostered a culture where people truly care for each other and the world around them.

Our reputation for software engineering and problem-solving was quickly established, and our customers started to view us more as their colleagues

We're much bigger now, helping organisations across the globe to navigate change in more ways than ever. But some things never change, like our commitment to our people, delighting our customers, and our technical ability

Some say we're leading a quiet revolution in digital, but if you ask us, we're just doing our bit to make the world a little bit better.



3,132 amazing
people



13 consecutive
years of growth



40 years
of innovation



£382m revenue
(FY2024)



13 locations
globally



FTSE 250 listed
tech company



Our core expertise

The digital revolution is already happening. We're leading from the front to help our clients **seize every digital opportunity** and be future-ready.



Cloud



Business Applications



Data and AI



Managed services



Engineering



Service and
experience design



Workday services
and products

- Test
- Audit
- Shield
- Deployment
- Optimisation
- Spark & Grow

Responsible AI

Responsible AI advisory and implementation services ensure the design and delivery of ethical, human-centric, secure, well-governed, and trustworthy AI systems. Our bespoke frameworks, workshops, and technical testing enable organisations to maximise value from advanced AI technologies while reducing risk across transformative AI initiatives.

Features

1. Responsible AI workshops to assess readiness across principles
2. Tailored frameworks for identifying, mitigating, and managing ethical risks
3. Fairness, bias mitigation, safety, and explainability testing and workshops
4. Human-centric design, people process, and culture conversations
5. EU AI Act, ISO 42001, NIST AI RMF alignment
6. Comprehensive risk identification, mitigation, and scenario planning
7. AI and Agentic AI governance frameworks, strategies, and compliance
8. AI third party risk evaluation and mitigation – audit and assurance
9. AI Security frameworks, red-teaming, and resilience planning against threats
10. Conducting ATRS, building model cards and other AI assurance services

Benefits

1. Minimise exposure to compliance failures, bias incidents, regulatory fines
2. Scalable AI deployment aligned with best practices and government standards
3. Trust in AI behaviour through testing, assurance, model correction
4. Accessible, inclusive, and user-aligned AI solutions
5. Embed governance structures and audit trails for accountability and oversight
6. Protect AI systems from threats and ensure business continuity
7. Ensure AI systems reflect ethical principles and organisational missions
8. Custom strategies to address sector and use-case specific risks
9. Build internal capability and confidence through training and best practices
10. Stay ahead of evolving regulations and standards with adaptable frameworks

Responsible AI

Service Overview

Artificial Intelligence is transforming every sector, but its adoption brings new risks and responsibilities. At Kainos, we believe that AI should be secure, trustworthy, ethical, compliant, and above all, human centric. Our Responsible AI services are designed to help organisations unlock the value of AI while ensuring it is developed and deployed safely, ethically, and in line with regulatory and societal expectations. We partner with clients to embed Responsible AI at every stage of the AI lifecycle, accelerating adoption, building trust, and safeguarding reputation.

Our Responsible AI Philosophy

Responsible AI at Kainos is the design and development of AI systems that:

- Uphold ethical principles

- Operate under strong governance with compliant action plans

- Foster trust through transparency and explainability

- Deliver measurable benefits to people and society

- Remain secure against external threats

We integrate these principles across five core themes: Ethics, Governance & Compliance, Trust, Human-Centric Design, and Security. Our approach is practical, actionable, and tailored to each client's context and maturity.

Service Catalogue

Kainos offers a suite of Responsible AI services, that can be contracted individually or incorporated as part of a larger AI delivery project. The typical stages of an AI delivery project are discovery, design, development, and deployment. We have made recommendations where relevant, below, on which stage is ideal for an offer, but can also incorporate our services at any stage if necessary.

Responsible AI Launchpad

We can provide a structured full-day workshop to evaluate your organisation's readiness across our 5 Responsible AI themes. During this workshop, we assess your current practices, identify gaps, and provide a clear, actionable roadmap for Responsible AI adoption. This is a high-level workshop, meant to capture general guidance and discussion, and to identify which areas of Responsible AI should be explored in further detail. The launchpad provides customers with a maturity assessment and high-level roadmap of actions and includes workshops such as an ethics and harms workshop to upskill the team in Responsible AI. This launchpad can be completed at an organisational or an individual AI use case level.

Ethical AI

For AI projects typically in the discovery stage, we offer a hands-on workshop to identify and evaluate the potential harms and benefits. This workshop will help stakeholders apply Responsible AI themes to the use case, gather diverse perspectives, and identify possible mitigations to take forward. It is a helpful starting point in upskilling your team on Responsible AI and ensuring Responsible AI is embedded from the beginning of a project.

For wider AI activity, we also offer a bespoke service to ensure AI is implemented ethically. We apply our AI Ethics Framework which covers mitigating harm, seeking benefits, responsibility, and respect for the humans behind the data to identify areas for development. As part of this service, we help you create and implement ethical guardrails, stakeholder engagement strategies, and sustainability impact assessments.

Responsible AI

Implementing Trust

For AI projects typically in the design, development, and deployment stages, we offer bespoke services in AI testing, trust and assurance. This includes AI testing on fairness, safety, explainability, bias mitigation, and transparency to ensure AI is functioning the way intended, or that AI behaviour is corrected in real-time. Our team is also able to apply wider practices using the Cross-Government AI Testing Framework as requested.

For wider AI activity, we also offer services in AI trust and assurance, such as developing transparency documentation and model cards, identifying and adding guardrails, and workshops on AI trust principles such as post-hoc explainability.

Government & Compliance

For organisations early in their AI Governance journey, we provide a bespoke service to help you set up your approach to Responsible AI Governance. This service includes applying our AI Governance Framework to conduct a Responsible AI maturity assessment, a gap analysis report with findings and recommendations to improve governance, a bespoke Responsible AI Strategy, a high-level AI policy, and a general approach to AI Governance for your organisation. This service will also set out a list of further actions to operationalise the strategy within your organisation. We offer these services as a package, but we can also create bespoke offerings based on your requirements.

For individual AI projects, we offer specific AI governance services, including sector-specific compliance assessments, industry

benchmarking, risk assessments and mitigations, AI registries and risk libraries, and regulatory mapping (e.g. EU AI Act, ISO 42001, NIST AI Risk Management Framework). We also offer a tailored Agentic AI service to navigate agentic activity and embed good governance. This work can take place at any stage of an AI lifecycle, but is recommended to take place during discovery, design, and development.

Security

For AI strategy and wider AI projects, we offer AI Security services to ensure your organisation has a strategy in place to deal with AI risks. As part of this service, we apply our AI Security Framework and tailor it to your organisation, to make recommendations on risk management, supply chain dependencies, secure development, and secure operations. This service is a bespoke offering that can include workshops, training, and reports based on a customer's needs.

We also offer practical AI Security technical services for AI projects in the design, development and deployment stages such as red-teaming, resilience planning, and secure pipeline design to ensure AI systems are protected from AI and non-AI based threats.

For wider organisations, we offer a bespoke third-party AI risk assessment and assurance service, to help you explore risks posed by vendors using AI and risk management techniques, as well as audit and assurance, and guidance on including this within procurement strategies.

Our Approach

Kainos' Responsible AI services are delivered through a structured, value-driven framework:

- **Discovery & Assessment**

We begin with collaborative workshops and interviews to understand your business challenges, strategic goals, and operational realities. We systematically identify and prioritise high-value AI use cases for your organisation, define clear business outcomes and KPIs, and assess your current responsible AI maturity.

- **Design**

We co-create solutions with stakeholders, ensuring AI augments human capabilities and optimises workflows. Our design process includes journey mapping, rapid prototyping, and embedding responsible AI governance from the outset. We address security, compliance, and ethical considerations at every stage.

- **Development & Deployment**

We use agile, iterative development to build and deploy AI solutions, leveraging reusable blueprints and accelerators for speed and reliability. We ensure seamless integration with existing systems and ongoing optimisation.

Tools & accelerators

Kainos is constantly developing new proprietary tools and accelerators to support Responsible AI adoption. Some of these tools include:

- AI Ethics Framework
- Responsible AI Launchpad
- AI Governance Framework
- AI Security Framework
- Agentic AI Governance Framework
- Responsible AI Dashboard
- Fairness metrics decision making
- Post-Hoc Explainability framework
- AI testing configurations

These assets enable rapid assessment, risk mitigation, and continuous monitoring, ensuring clients stay ahead of regulatory and market expectations. They also provide guidance on AI and Agentic AI testing in the context of Responsible AI considerations such as ethics, bias, and fairness to ensure that AI works as intended or if not, that AI behaviour is corrected in real-time. We also offer a Responsible AI Dashboard tool which gives a single-pane-of-glass view of assurance to assist with decision making.

These tools are designed to help with Responsible AI decision making so that AI can be scaled with confidence, trust, and assurance.

Case Studies

Kainos' Responsible AI services have delivered measurable outcomes for clients across sectors, such as:

- **Justice sector:** Developed an AI strategy focused on ethics, trust, and governance, resulting in a clear roadmap, improved decision-making, and enhanced justice equity.
- **Research sector:** Enabled safe, responsible AI experimentation in highly secure environments, integrating ethics, human factors, and data science.
- **Local government sector:** Designed and implemented an AI-powered chatbot for financial support, developed in line with Responsible AI standards, improving accessibility and freeing up frontline staff.
- **Financial sector:** Embedded ethical analysis and transparency documentation into AI proof of concepts, driving responsible innovation.

These engagements demonstrate our ability to deliver practical, scalable Responsible AI solutions that build trust, reduce risk, and accelerate adoption.

Responsible AI

Why Kainos

Proven Blueprint: Deep experience in public sector and regulated industries, with a reputation for ethical delivery and best practice development.

Technical Depth: Unique technical capabilities to turn Responsible AI principles into practice, including advanced testing, explainability, and security.

Thought Leadership: Active participation in shaping international policy and standards, with strong connections to industry, academia, and government.

Comprehensive Support: End-to-end services, from strategy and assessment to implementation and continuous improvement.

Kainos' Responsible AI services provide a holistic, actionable pathway to safe, ethical, and effective AI adoption. By embedding Responsible AI at every stage of the AI lifecycle, we help organisations accelerate innovation, build trust, and achieve measurable business value - while safeguarding people, reputation, and compliance.



Where we are

The Americas

Buenos Aires
Indianapolis
Nova Scotia
Toronto

UK and Ireland

London
Birmingham
Belfast
Derry
Dublin

Central Europe

Gdansk
Helsinki
Paris
Wommelgem



What makes us stand out



Our people



True collaboration



Innovation



Digital know-how



Common goals



Trust



Digital Services

We are **engineers**, specialists. We **overcome big challenges** for our clients, through delivery of **intelligent** digital services that use the best in **talent and technology**.



55 million UK citizens

positively impacted by
services we've delivered



**892
customers**



13 locations

And growing



99% of customers

rate our service as good,
great or excellent




Key partners

Working closely with
Microsoft and AWS to
deliver results



Award-winning

90+ industry awards



Workday

Kainos is a Workday **phase one prime status partner** in all European, North American and Asia Pacific markets. We have worked with many public and private sector customers of various sizes and scale in over **40 countries worldwide**.



Partner Since 2011

We are the only certified Workday Services, Technology and Extend partner globally.



Workday Services

Full range provider from pre-deployment consulting, implementation and support.



650+ Customers

We have deployed to over 650 customers and are a proud Workday customer also, using HCM, Financials, Recruitment, Planning and PSA to run our business.



Workday Public Sector

We have worked with many UK Public Sector customers with their initial deployment, post-deployment support and adoption of new modules.



Kainos Smart

Our product is deployed to 500+ customers globally saving them 1000s of testing hours.



900+ Consultants

Hold over 2,700+ Workday certifications covering all areas of the platform.

65 million

citizens positively impacted
by services **built by Kainos**
for the UK Government



Every driver & every car on the road
Every piece of work to fix roads
Every court case & every prisoner
Every home helped with help-to-buy
Every passport & renewal application
Every COVID test sent out
Every COVID test result
Every user of the NHS App

Every voter registered
Every company opened & closed
Every UK civil servant trained
Every piece of produce grown in the UK
Every thing imported & exported to the EU
Every benefit fraud prevented
Every body's pension details

Our Social Values and commitment to HM Government strategy



Education & Capability Building

We empower people through technical and behavioral learning programs, academies, and outreach projects, fostering skills development and inspiring future technology leaders across communities and organisations.

Social Responsibility

Tackling workforce inequality and ensuring equal access to technical and vocational education, through our Academies, Camps and Earn as you Learn schemes.

Innovation & Future Readiness

Driving continuous improvement through technology

Health, Wellbeing & Accessibility

We support colleagues' physical, emotional, and career wellbeing through initiatives, guidance, and inclusive design, ensuring equal access, experiences, and compliance for all users with diverse needs.

Climate Action & Sustainability

With a comprehensive sustainability strategy underway, we will be a Carbon Net Zero company by 2030.

Equality, Diversity & Inclusion

Maximising service experience for citizens and civil servants by delivering projects with diverse teams, underpinned by our BAME, LGBTQIA+ and female communities.



The Code of Ethics

We are guided by our 6 ethical principles.



Wellbeing

We protect the wellbeing of our staff, customers and communities.



Equality

We improve access and inclusion and minimise bias.



Transparency

Our decisions are traceable and accountable.



Integrity

We hold ourselves and others to ethical standards.



Environment

We act responsibly towards the Earth and its resources.



Initiative

We take initiative to deliver social value and positive impacts.

BCDR and Exit Plan

Business Continuity and Disaster Recovery (BCDR)

A Business Continuity Plan can be provided if required. This shall set out the arrangements to be invoked in the event of an actual or perceived threat to business continuity, to ensure continued operation of the system and continuity of the services provided by Kainos pursuant to the Prime Agreement and shall include: the alternative processes, options and responsibilities that may be adopted in the event of a failure or disruption to the system and/or services provided by Kainos pursuant to the Prime Agreement; and the steps to be taken by Kainos upon resumption of the system and services provided by Kainos pursuant to the Prime Agreement in order to address any prevailing effect of the failure or disruption including a root cause analysis of the failure or disruption.

Exit Plan

An Exit Plan can be provided if required to detail the steps that would be carried out to ensure smooth transition of Kainos services to a new supplier. The steps outlined in the Exit Plan will help mitigate against any disruption to the service during the transition period. It is assumed that any new supplier will themselves have a procedure they wish to follow during the transition period, and as such the steps in this Exit Plan will serve as a checklist for the new supplier to ensure all key areas of the transition have been covered. The Exit Plan will therefore be subject to refinement should it be exercised.

Commercial Statement

Confidentiality and copyright

© Kainos Software Limited 2026 ("Kainos")

The contents of this document are commercial and confidential in nature and the copyright of Kainos. This document must not be reproduced (in whole or in part) save in connection with the purpose for which it was issued and must not be shared with Generative AI systems without express consent.

Trademarks

Kainos® is a registered trademark of Kainos Software Limited. All rights reserved. You may not delete or change or modify any copyright or trademark notice.

Caveats

Kainos has used all reasonable endeavours to ensure that the contents of this document are accurate but is not responsible for any errors or omissions.

All information provided prior to execution of a contract is provided 'as is' and 'subject to contract' without warranty of any kind.

This document does not constitute an offer from Kainos. In the event that the parties elect to work together, they will only be contractually bound to each other upon signature of a contract.

Corporate information

"Kainos" is the trading name of the Kainos group of companies, further information on which can be found here: <https://www.kainos.com/corporate-information/>.

Kainos Software Limited/ Kainos WorkSmart Limited, will be the contracting entity under this tender / for the provision of the service and may be assisted from time to time by other Kainos group companies.

Freedom of information

Kainos considers that the following information provided in this document is exempt from disclosure under the Freedom of Information Act 2000 (FOI):

The CVs of staff qualify under the "Personal Information Exemption (s.40)" of the Freedom of Information Act and are exempt from disclosure under the Data Protection Act 1998. The Period for which this information should be confidential is the lifetime of the Data Subject.

Rate and pricing information is confidential and commercially sensitive and covered by the 'Commercial Interests' exemption (s.43) of the FOI, as the release of this information is likely to prejudice the commercial interests of Kainos and is likely to adversely affect its (and the Customer's) future negotiating position. The period that this information should be confidential for should be 5 years.