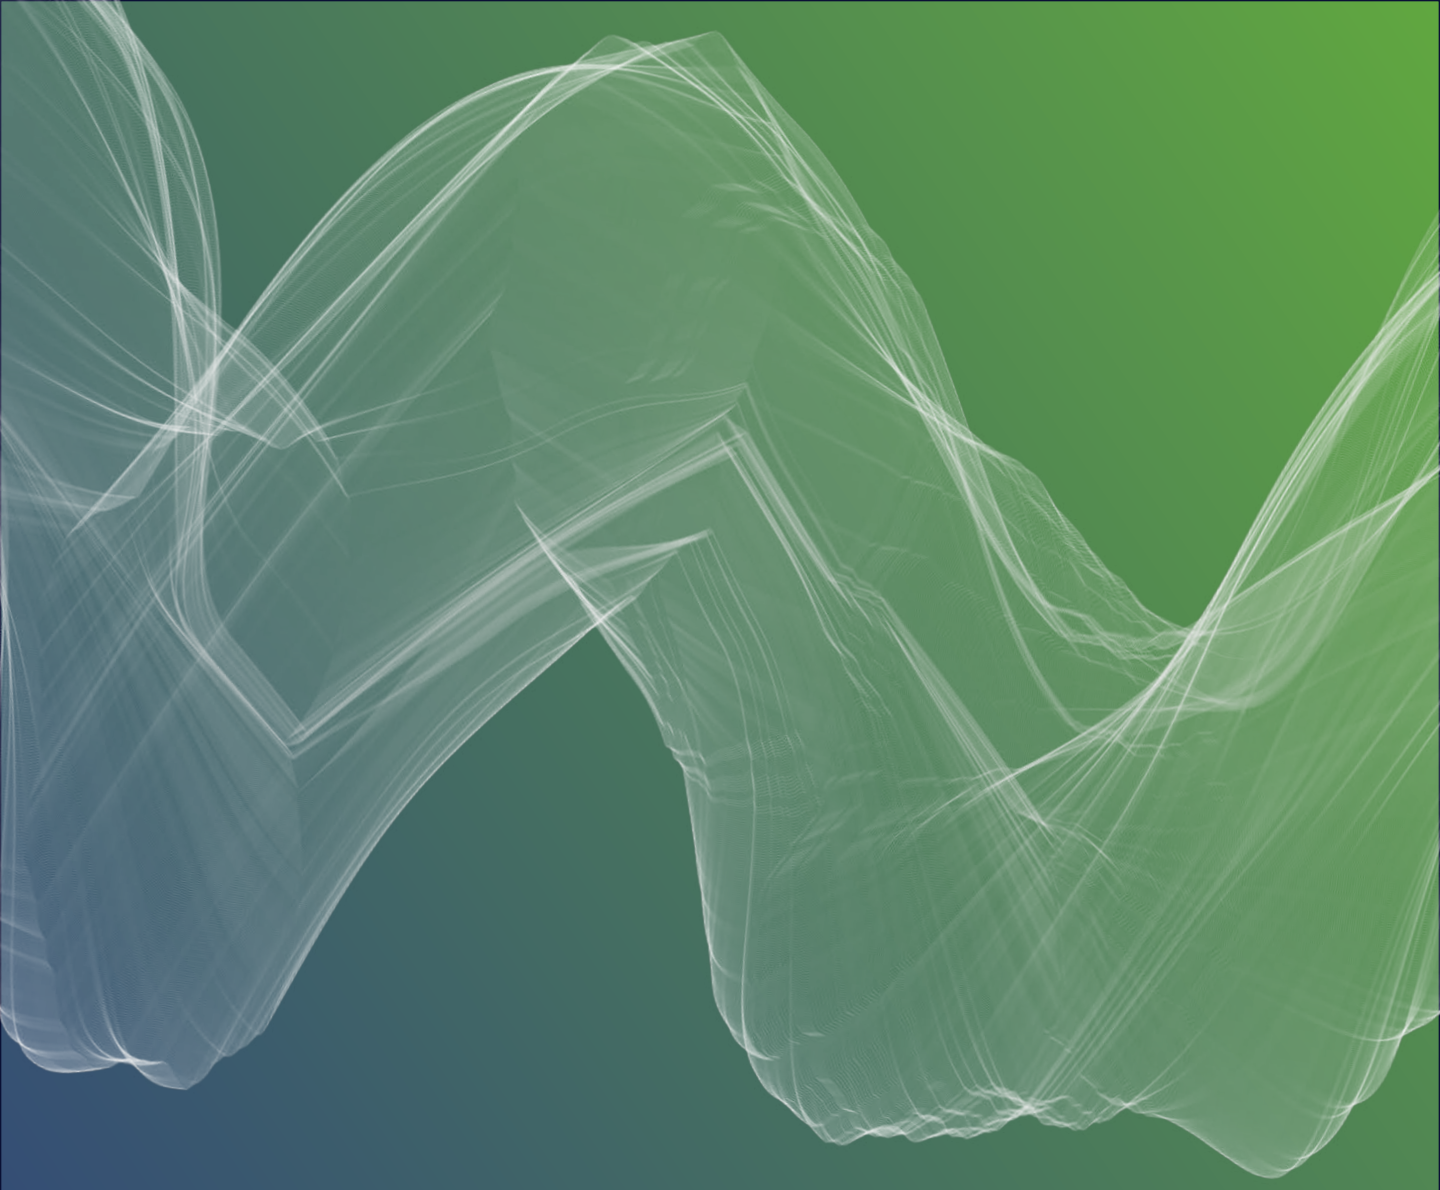

kainos®

Cyber and AI Security Discovery

Service Definition Document

G-Cloud 15

Status: Definitive
Security Commercial – In Confidence
Version no.1/0

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

Cyber & AI Security Discovery

Cyber/AI Security Discovery ensures customers are provided with complete visibility of all AI, data, and IT assets, real-time validation of security control effectiveness, and quantified risk prioritisation across both AI and traditional systems - so you can spot control gaps, respond fast, and execute a clear improvement roadmap for governance, protection, and continuous monitoring.

Features

1. **Asset/Data Discovery,** Asset inventories, data flow diagrams, classification tags.
2. **Current-state Controls Assessment**, Gap analysis , maturity scores, prioritised action lists.
3. **Risk Assessment Execution,** Risk register, risk treatment recommendations, exposure dashboards.
4. **Third-Party/Supply Chain Mapping,** Third-party inventories, risk ratings, supplier assurance checklists.

Benefits

1. Eliminate shadow AI blindspots and reduce hidden asset risk exposure
2. Quantified risks based on real materialisable threats, not theoretical checklists
3. Board-ready visibility into assets, controls, risks, investment priorities, governance clarity
4. Reduced fatigue via audit-ready artefacts, evidence links, compliance proof mechanisms
5. Supplier risk acceleration with contractual control clarity through AI-specific diligence.
6. Governance foundation feeding directly into GOVERN, PROTECT, DETECT workstream prioritisation
7. Cost avoidance from risk detection, smarter remediation, reduced rework downstream
8. Standards alignment for EU AI Act, NIS2, GDPR, sector-specific-regulatory frameworks
9. Improved insurability demonstrable oversight of third-party AI dependencies, control coverage
10. Faster approvals via objective baseline recognised by boards/auditors/regulators/stakeholders

Cyber & AI Security Discovery

Quantification, and Control Validation
Artificial intelligence is accelerating across government and regulated enterprises, creating attack surfaces that traditional cyber inventories cannot see. Organisations deploy foundation models, fine-tuned generative AI, agentic workflows, and third-party AI services at pace, yet most remain blind to where these assets live, what risks they pose, and whether their existing security defences actually work. Shadow AI and Copilot sprawl make the "unknown unknowns" a board-level risk. Simultaneously, emerging regulatory obligations—the EU AI Act, NIS2, ETSI technical standards, and GDPR—demand demonstrable discovery, quantified risk assessment, and governance baselines. Organisations need clarity fast: what do we have? Where are the real risks? How do our controls actually perform?

Cyber/AI Security Discovery is the answer.
It ensures complete visibility of all AI, data, and IT assets, real-time validation of security control effectiveness, and quantified risk prioritisation across both AI and traditional systems. The service delivers rapid visibility, quantified risk, and board-ready improvement roadmaps in just 4–6 weeks—so you can spot control gaps, respond fast, and execute a clear improvement roadmap for governance, protection, and continuous monitoring.

What Cyber/AI Security Discovery Delivers
Cyber/AI Security Discovery transforms uncertainty into actionable intelligence through four interconnected service components:

1. Asset & Data Discovery
We build a complete, AI-aware inventory of everything your organisation relies on—models (foundation, fine-tuned, locally hosted, cloud-hosted), datasets, prompts, agents, AI applications, MLOps pipelines, cloud platforms (Azure, OpenAI, AWS Bedrock, etc.), infrastructure, data flows, and conventional IT assets. Using the Kainos AI Security Assurance Framework (AISAF) discovery-readiness toolkit and, where relevant, our Multi-Cloud AI Threats Discovery Engine, we integrate seamlessly with your existing CMDB (ServiceNow) to enrich configuration items and link to business-technical services. The result is a single source of truth for governance, monitoring, and change control—eliminating shadow AI and preventing blindspots that defeat control design.

2. Third-Party & Supply Chain Mapping
AI supply chains are complex: models, datasets, and agents flow from multiple vendors, partners, and API providers. Systemic risk propagates rapidly, and insurers and regulators increasingly expect demonstrable oversight. We catalogue AI-related suppliers and perform AI-specific due diligence—evaluating model provenance, data handling, security attestations, contractual controls, and risk exposure. We uplift existing third-party risk management (TPRM) frameworks with AI-specific categories, provide a supplier assurance checklist, and create dependency mapping visuals. Where access allows, we support shadow validation and independent testing alongside vendors, ensuring you know exactly which external dependencies matter and where controls are missing.

Cyber & AI Security Discovery

3. Risk Assessment Execution
Checklist audits miss AI-specific failure modes. Organisations need business-context risk quantification. We apply a proven adversarial threat modelling methodology—incorporating OWASP Top 10 for LLMs, MITRE ATLAS, and OWASP Agentic Security frameworks—combined with organisation-specific impact factors: decision criticality, data sensitivity, service integration level, and supply-chain exposure. Risks are scored for likelihood and impact, capturing inherent and residual risk, and prioritised based on what can actually materialise in your environment. Outputs are captured in a structured risk register with defined owners, treatment timelines, and treatment options. We produce a risk exposure dashboard that enables executives to see quantified, prioritised exposure and make investment decisions grounded in real threat intelligence, not tick-box compliance.

4. Security Posture Assessment
Your existing controls need objective benchmarking. We assess current controls and practices against a harmonised catalogue mapped to NIST CSF 2.0 (including the GOVERN function), NCSC CAF, ISO 27001, ISO 42001, and ETSI TS 104 223 (the UK AI Cyber Security Code of Practice). We score maturity across each control domain, conduct gap analysis, and produce a compliance heatmap that cross-walks your findings against GDPR, EU AI Act, NIS2, and sector-specific frameworks (health, finance, transport, etc.). These findings integrate directly with your GRC platform (e.g., OneTrust) and feed a prioritised improvement roadmap that sets up subsequent PROTECT and DETECT delivery.

Methodology & Delivery Approach
Cyber/AI Security Discovery follows a structured five-phase delivery model:

Week 0–1: Mobilise & Scope
We confirm objectives, inscope business units and systems, connect to your CMDB and GRC platform, and agree on evidence collection plans and data handling protocols.

Week 1–3: Discover
We execute AISAF discovery across four domains—Risk Management, Supply Chain, Software Development Lifecycle (SDLC), and Secure Operations. We build a consolidated AI-and-IT asset inventory, preliminary data flow diagrams, and collect supplier inventory data.

Week 2–4: Assess Risk
We conduct workshop-led adversarial threat modelling, identify AI-specific and conventional security risks, and quantify them in a structured risk register with treatment options drafted.

Week 3–4: Benchmark Controls
We score control maturity against NIST CSF 2.0, CAF, and ETSI standards, produce a compliance heatmap, and analyse control gaps.

Week 4–6: Synthesis & Playback
We deliver a consolidated executive report, risk exposure dashboard, and executive briefing with a prioritised improvement roadmap and investment cases that feed directly into GOVERN, PROTECT, and DETECT workstreams.

Typical engagement duration: 4–6 weeks, conducted with a blended Kainos team (AI Security Lead, Cyber Security Engineers, Data/AI Subject Matter Expert, part-time Delivery Manager) and client co-delivery (security, data, and service owners embedded), ensuring capability transfer from day one.

Cyber & AI Security Discovery

Key Deliverables

- **AI-aware Asset Inventory:** Unified catalogue of AI-and-IT assets, integrated with CMDB, serving as the single source of truth for governance, monitoring, and change control.
- **Data Flow Diagrams:** Authoritative diagrams of training-inference data paths, model endpoints, and integrations underpinning threat modelling, privacy, and access reviews.
- **Third-Party Inventory & Risk Ratings:** Supplier list with AI-specific due diligence scoring and assurance checklist, informing TPRM decisions, contract clauses, and targeted testing.
- **Risk Register & Treatment Plan:** Structured register with risk ID, description, likelihood, impact, inherent and residual scoring, owner, and treatment, prioritising remediation and feeding risk dashboards and board reporting.
- **Control Gap Analysis & Maturity Scores:** Scoring against NIST, CAF, ISO, and ETSI standards, evidence-linked gaps, and a sequenced improvement roadmap with measurable targets (Governance Maturity Index, Assurance Coverage Ratio).
- **Compliance Heatmap:** Crosswalk to GDPR, EU AI Act, NIS2, and sector frameworks, providing audit-ready proof of readiness.
- **Executive Report & Roadmap:** Synthesised findings, quantified exposure, investment-prioritised improvement plan, and sequenced delivery roadmap.

Governance & Standards Alignment

Cyber/AI Security Discovery assessments align to NIST CSF 2.0 (including GOVERN), NCSC CAF, ISO 27001, ISO 42001, and ETSI TS 104 223—standards core to UK, European, and international compliance. The compliance heatmap maps findings against GDPR, UK GDPR, EU AI Act, and sector-specific frameworks (health, financial services, transport, etc.). All outputs are audit-ready, evidence-linked, and suitable for regulatory and internal audit programmes. We leverage Kainos' thought leadership—we authored the DSIT AI Cyber Security Code of Practice Implementation Guide and co-lead OWASP Top 10 for LLMs and OWASP Agentic Security—ensuring guidance is adoption-led, practical, and grounded in real-world delivery.

Why Organisations Need This Now

Government and regulated sectors face escalating AI-driven risks: prompt injection, data leakage, adversarial manipulation, and supply-chain compromise. Boards expect quantified exposure and defensible prioritisation. The convergence of standards (UK AI Code of Practice becoming ETSI TS 104 223) across EU and ISO ecosystems makes standards-based discovery and risk mapping an immediate prerequisite for audit readiness and conformity schemes. Cyber/AI Security Discovery closes the gap between AI ambition and secure, governed execution.

Cyber & AI Security Discovery

Outcomes & Impact

Customers typically gain:

- **Complete visibility** of all AI, data, and IT assets, eliminating shadow AI and blindspots.
- **Real-time validation** of security control effectiveness, moving beyond checklist compliance to proven defence capability.
- **Quantified, materialisable risks** grounded in business context, not theoretical assessments, reducing report fatigue and accelerating investment decisions.
- **Third-party exposure insights** with contractual control clarity and supplier risk ratings.
- **Board-ready improvement roadmaps** that feed governance, protection, and detection workstreams with confidence.
- **Cost avoidance** from early risk detection, accelerated approvals via audit-ready evidence, reduced rework in build-and-operate phases, and improved insurability through demonstrable third-party oversight.
- **Audit readiness** with evidence-linked artefacts, control mappings, and compliance proof suitable for regulators and assurance programmes.

Scalable Engagement Options

- **Rapid Assessment (4–6 weeks):** Single business unit or service, focused on surfacing AI-and-IT assets, top risks, and priority gaps—ideal pre-pilot or pre-go-live.
- **Deep-Dive Assessment:** Enterprise-wide discovery, supplier inventory, control benchmarking, executive reporting, and investment roadmap for multi-site, multi-platform organisations.
- **Phased Discovery:** Rolling approach by platform, business unit, or risk tier, building the baseline iteratively without disrupting delivery.

Engagement & Next Steps

To initiate a Cyber/AI Security Discovery assessment, schedule a discovery call covering your AI usage and ambitions, known-and-unknown AI data assets, third-party AI dependencies, current governance and GRC platforms, and compliance drivers (EU AI Act, NIS2, GDPR). We will propose a right-sized scope (Rapid, Deep-Dive, or Phased), define evidence needs, and outline timescales and team composition. Contact your Kainos account team or the AI Security inbox to schedule your scoping session.

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

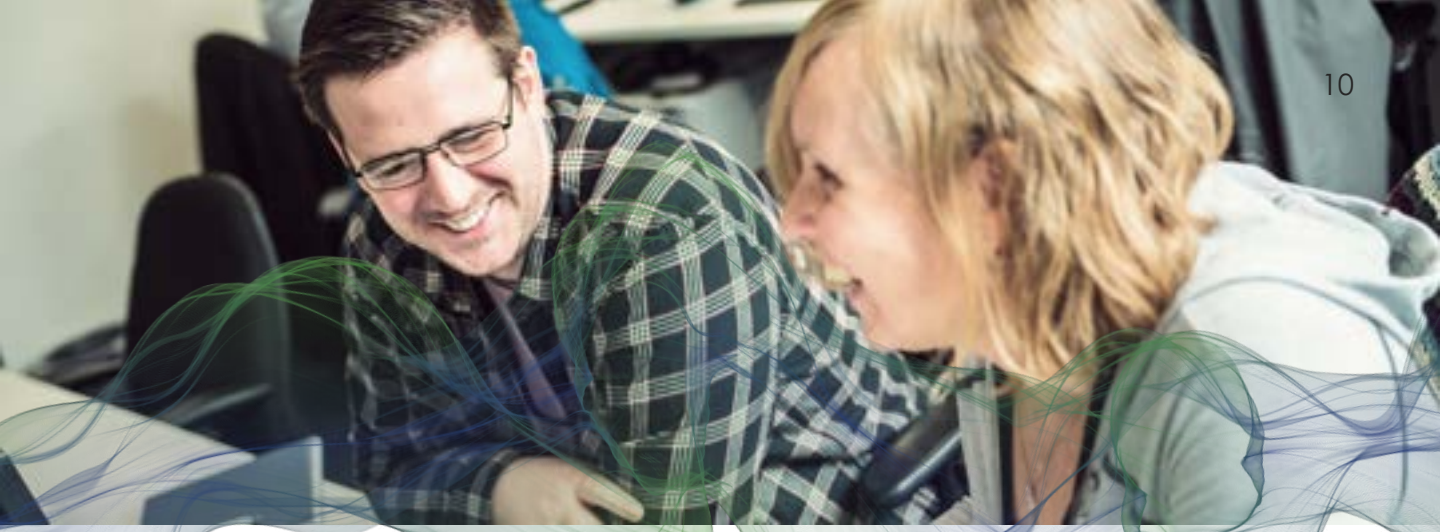

What makes us stand out

- 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**.

NHS
Ministry of Defence
Department for Environment Food & Rural Affairs
Department for Transport
UK Health Security Agency
Home Office
Driver & Vehicle Standards Agency
Ministry of Justice
HM Revenue & Customs
Foreign & Commonwealth Office

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

kainos®

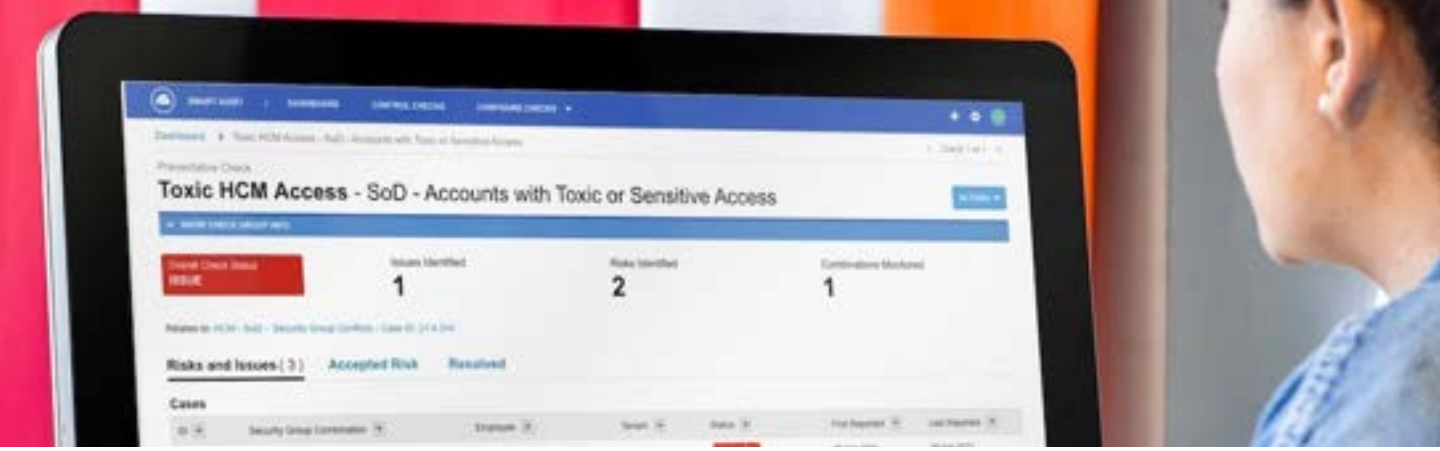

Workday

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**.

aggreko | KION GROUP | Magellan HEALTH | St. Luke's Health | The National Archives

Crown Commercial Service | Department for Education | NETFLIX | UKRI Innovate UK | Nasdaq

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

kainos®

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

Trademarks

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.