SERVICE DEFINITION DOCUMENT

Service Name: AI Security Assurance & Safety: LLM, Generative AI & Cloud Governance

Provider: Secure Equinox Limited

G-Cloud 15

Reference: Jan 2026

---

1. Service Overview

Secure Equinox Limited provides specialist **AI Security Assurance and Governance** services. As the UK public sector accelerates its adoption of Generative AI and Large Language Models (LLMs), we act as your strategic partner to ensure these technologies are adopted safely, legally, and securely.

We do not just "block" AI; we enable innovation by translating complex AI risks (such as Prompt Injection, Hallucination, and Data Leakage) into clear, actionable governance frameworks aligned with the NCSC Guidelines for Secure AI System Development and the AI Safety Institute.

2. Key Service Features

- Generative AI Risk Assessment and Safety Governance Frameworks.
- Microsoft 365 Copilot security configuration and data loss prevention (DLP) policy design
- "Secure by Design" Architecture for AI and LLM adoption (AWS/Azure).
- AI Policy Development and Acceptable Use Policy (AUP) creation.
- NCSC-aligned AI Security Principles and Guidelines Implementation.
- Prevention of Prompt Injection and Data Poisoning vulnerabilities.
- AI Model Assurance and Supply Chain Risk Management (SCRM).
- Compliance with AI Safety Institute and Global Standards (ISO 42001).
- EU AI Act readiness assessment for organizations with EU data flows.
- Data Privacy Impact Assessments (DPIA) for specific AI use-cases.
- Strategic AI Security Leadership and Virtual CISO Services.

3. Key Service Benefits

- **Safe Innovation:** Enable your organisation to use AI tools (e.g., Copilot, ChatGPT Enterprise) without compromising security.
- **Regulatory Compliance:** Ensure alignment with emerging UK Government AI frameworks and GDPR requirements.
- **Reputation Protection:** Mitigate risks of bias, toxicity, and "hallucinations" in public-facing AI services.
- **Data Sovereignty:** Architecture designs that ensure sensitive government data does not leak into public models.
- **Clear Governance:** We define *who* is responsible for AI risk, moving you from "Shadow AI" to managed capability.

4. Key Deliverables

All engagements produce formal, auditable documentation including (but not limited to):

Core Risk & Governance:

- AI Usage Inventory mapping sanctioned and unsanctioned AI tools with risk ratings
- AI Risk Register with quantified likelihood/impact scores aligned to your risk appetite
- AI Governance Framework defining policies, processes, and decision-making authorities
- AI Governance RACI Matrix with clear accountability for AI decisions and risk ownership
- AI Acceptable Use Policy and guardrails for safe staff usage

Compliance & Assurance:

- Data Protection Impact Assessment (DPIA) for specific AI use-cases
- AI Impact Assessment evaluating fairness, bias, and consequences on individuals
- Algorithmic Transparency Recording Standard (ATRS) documentation for public-facing AI systems
- AI Supply Chain Risk Assessment for third-party models, APIs, and components

Technical Documentation:

- AI Model Documentation including data provenance, scope, limitations, and failure modes
- Cloud Architecture Diagrams showing secure AI deployment (Azure/AWS)
- Security Configuration Guides for AI platforms (e.g., Microsoft Copilot, ChatGPT Enterprise)
- AI System Monitoring Framework with drift detection and logging procedures

Operational Readiness:

- AI-specific Incident Response Plan covering prompt injection, poisoning, and hallucination events
- AI Security Awareness Training materials for staff on safe usage and reporting

Note: Deliverable selection is tailored to engagement scope and objectives. Not all deliverables are required for every engagement. During onboarding, we agree a specific deliverable list matching your accreditation requirements, organizational standards, and budget constraints.

5. Our Approach & Methodology

We utilize a risk-centric approach tailored to the unique challenges of probabilistic AI models:

- **Discovery:** We map your current AI usage (sanctioned and unsanctioned) to understand the "Shadow AI" landscape.
- **Assessment:** We utilize NCSC Risk Management principles and ISO 42001 (AI Management System) standards to assess the security of your proposed AI solution.
- **Policy & Governance:** We write the "Rules of the Road"—creating Acceptable Use Policies and Guardrails that tell your staff *how* to use AI safely.
- **Architecture:** We review your Cloud (AWS/Azure) environment to ensure the "hosting" of the AI model is secure, even if the model itself is a black box.

6. Onboarding and Offboarding

- **Onboarding:** Engagement begins with a structured "AI Safety Workshop" to identify your highest-priority AI use cases.
- **Offboarding:** We transfer all Risk Assessments, Policies, and Architecture designs to your internal teams. We ensure your permanent staff understand how to monitor the AI system moving forward.

7. Service Constraints

- Consultants are UK-based.
- We focus on Governance and Assurance. We do not provide software development or data science engineering services.
- Access to technical documentation for third-party models may be required for deep-dive assurance.

8. Pricing

Pricing is based on our G-Cloud SFIA Rate Card.

- **Standard Rate:** Our published rates apply across all relevant SFIA categories (Strategy & Architecture, Business Change, Solution Implementation, etc.), reflecting the cross-functional nature of AI Governance and Security Assurance.
- **Discounts:** Educational or volume discounts are not generally applicable due to the specialist nature of our resources. However, preferential rates for significant, strategic engagements (e.g., multi-year programmes) may be discussed at the Call-Off stage.

[End of Document]