



AI 'in a Safe Box' Secure AI Adoption for Sensitive Data (LLMs, RAG, Guardrails)

G-Cloud 15 – Lot 2b (Cloud Software)





Table of Contents

1. Service overview	3
2. Intended users	3
3. Core capabilities	3
3.1 Secure AI platform and lifecycle management	3
3.2 Large language models and agentic workflows	3
3.3 Retrieval-augmented generation (RAG)	3
3.4 Guardrails and human oversight	3
3.5 Evaluation, testing, and assurance	4
4. Trustworthy AI and governance	4
5. Deployment model	4
6. Onboarding and configuration	4
7. Integration and APIs	4
8. Security and data protection	4
9. Support and service management	4
10. Service constraints	5
11. Offboarding and data export	5
12. Benefits to buyers	5
13. Alignment to G-Cloud Lot 2b	5

Secure AI Adoption for Sensitive Data (LLMs, RAG, Guardrails)

"AI in a Box" – Service Definition

G-Cloud 15 – Lot 2b (Cloud Software)

1. Service overview

Adappt provides a deployable cloud software platform enabling the secure adoption of large language models (LLMs) and agentic AI for sensitive data. The service is delivered as an "AI in a box" solution hosted in a secure, buyer-controlled private cloud environment. It supports retrieval-augmented generation (RAG), policy guardrails, evaluation, and audit controls to enable governed, trustworthy AI use.

2. Intended users

- Public sector organisations handling sensitive or regulated data
- Policy, regulatory, and compliance teams
- Digital and data teams deploying AI-enabled services
- Security, assurance, and risk management functions
- Operational teams adopting AI for decision support or automation

3. Core capabilities

3.1 Secure AI platform and lifecycle management

The platform provides core AI capabilities to support the full lifecycle of AI workloads, including configuration, deployment and controlled updates. AI components are versioned and governed to support repeatability, auditability, and controlled change.

3.2 Large language models and agentic workflows

The service supports integration and operation of large language models and agentic workflows within a secure environment. AI behaviour is constrained through configurable prompts, rules, and policies aligned to buyer governance requirements.

3.3 Retrieval-augmented generation (RAG)

RAG capabilities enable LLMs to retrieve information from approved data sources at runtime. Access controls, permissions, and traceability ensure AI outputs are grounded in authorised data and can be audited.

3.4 Guardrails and human oversight

Policy-based guardrails control prompt usage, outputs, and actions. Human-in-the-loop approval steps can be configured where required, supporting high-assurance use cases and reducing operational risk.



3.5 Evaluation, testing, and assurance

The platform includes evaluation and testing capabilities to assess AI behaviour, accuracy, bias, and safety. These capabilities support continuous assurance and informed decision-making prior to and during production use.

4. Trustworthy AI and governance

The service is designed to support trustworthy and responsible AI adoption. Governance features include role-based access control, audit logging, traceability of actions and outputs, transparency of automated decisions, and alignment with organisational policies and standards.

5. Deployment model

- Private cloud (buyer-controlled environment)
- Hybrid cloud (where agreed)

6. Onboarding and configuration

Standard onboarding is included to configure AI environments, RAG components, guardrails, access controls, and integrations. Onboarding is delivered remotely and supports effective use of the platform. Bespoke software development is not included.

7. Integration and APIs

The service provides APIs to integrate with existing data sources, document repositories, applications, identity providers, and key management services. APIs enable secure ingestion, retrieval, and interaction with AI capabilities.

8. Security and data protection

Security controls include encryption in transit (TLS 1.2 or higher), encryption at rest where applicable, integration with buyer-managed key management services, role-based access control, segregation of customer data, and logging and monitoring of access and activity. Buyer data remains under buyer control at all times.

9. Support and service management

Support is provided via email or online ticketing during Working Hours (09:00–17:30 UK, Monday to Friday). Requests are prioritised by severity. Enhanced support options, including extended hours or 24/7/365 support for production AI services, are available by prior agreement at additional cost.

10. Service constraints

Service capabilities are influenced by selected LLM providers, model limitations, data quality, and buyer governance requirements. Performance and behaviour depend on the deployed cloud environment and retrieval sources. Planned maintenance may be required for platform updates and will be communicated in advance.

11. Offboarding and data export

During an agreed offboarding period, buyers can export their data using available tools or APIs in commonly used formats. Following offboarding, remaining data is securely deleted in line with contractual and security requirements.

12. Benefits to buyers

- Secure adoption of AI for sensitive or regulated data
- Reduced risk through governed and auditable AI use
- Improved accuracy and relevance through RAG
- Greater trust and transparency in AI-assisted decisions
- Buyer control over data, environments, and policies
- Faster transition from pilot to production AI services

13. Alignment to G-Cloud Lot 2b

This service is provided as cloud software under Lot 2b. Configuration and onboarding activities are limited to enabling effective use of the platform and do not include bespoke development or managed operational services. Optional implementation or adoption support may be procured separately under Lot 3 where required.