

Ai Rag Implementation

G-Cloud

January 2026

Table of Contents

1	Service Overview
2	Our Approach
3	Getting Started
4	Your Responsibilities
5	Security & Compliance
6	Support & Exit
7	Why Appvia
8	Contact Us

AI and RAG Implementation

End-to-end implementation of generative AI solutions including Retrieval Augmented Generation (RAG), knowledge base assistants, and document processing with LLMOps pipelines.

1 Service Overview

Appvia delivers production-ready RAG solutions that enable natural language access to your enterprise knowledge. We implement using Amazon Bedrock, Azure OpenAI, and Google Vertex AI with comprehensive evaluation pipelines.

What you get:

- **RAG architecture** — Enterprise knowledge base design and implementation
- **Foundation model selection** — Structured evaluation to choose the right model
- **RAGAS evaluation** — Faithfulness, relevancy, and precision metrics
- **Vector database** — Implementation and retrieval optimisation
- **Prompt engineering** — Systematic testing and versioning
- **Guardrails** — Content filtering and safety evaluation
- **LLMOps observability** — Production monitoring dashboards
- **CI/CD pipelines** — Automated model and prompt deployment

We deliver measurable accuracy with reduced hallucinations through evaluation pipelines.

2 Our Approach

We follow a structured approach with evaluation at every stage to ensure production-ready AI.

Implementation phases:

Phase	Activities
-------	------------

Design	RAG architecture, model selection, retrieval strategy
Build	Knowledge base ingestion, vector store setup, prompt development
Evaluate	RAGAS pipeline, faithfulness testing, retrieval optimisation
Deploy	Production deployment, monitoring, CI/CD automation

Our evaluation approach: - RAGAS metrics for faithfulness, answer relevancy, context precision - LLM-as-a-judge for automated quality assessment - Systematic prompt testing with version control - Continuous evaluation in production

We focus on transparent AI with source citations and audit trails.

3 Getting Started

RAG implementations typically take 8-12 weeks depending on complexity:

Phase 1: Discovery We assess your knowledge sources, define use cases, and design RAG architecture. We evaluate foundation model options.

Phase 2: Build Knowledge base ingestion, vector database setup, prompt engineering, and initial RAG pipeline development.

Phase 3: Evaluate & Optimise RAGAS evaluation pipeline, retrieval tuning, faithfulness optimisation, and guardrails implementation.

Phase 4: Deploy Production deployment, LLMOps monitoring, CI/CD setup, and knowledge transfer.

Prerequisites: - Defined knowledge sources (documents, databases, APIs) - Cloud platform access (AWS, Azure, or GCP) - Sample queries and expected outcomes for evaluation - Stakeholders for use case validation

4 Your Responsibilities

To ensure successful service delivery, you will need to:

- Provide timely access to required resources, stakeholders, and information

- Make appropriate subject matter experts available for design and integration discussions
- Ensure business users are available for acceptance testing as agreed
- Procure necessary software licenses and subscriptions
- Empower your project team to make decisions in a timely manner
- Review and approve deliverables based on agreed acceptance criteria
- Provide access to knowledge sources and documents
- Define sample queries and expected answer quality
- Ensure cloud platform access for RAG infrastructure
- Participate in evaluation and quality validation
- Approve guardrails and content filtering policies

5 Security & Compliance

This service supports workloads classified up to **Official-Sensitive** where appropriate controls are in place.

Our security credentials:

- **ISO 27001:2013** and **ISO 9001:2015** certified
- **Cyber Essentials Plus** certified
- **75%+ of our team hold SC clearance**
- All access to customer environments is logged and auditable
- Data remains within UK regions unless explicitly agreed

We operate in alignment with GDS, NCSC, and CDDO security principles.

6 Support & Exit

Ongoing Support

- Dedicated Service Delivery Manager as your primary contact
- Regular service reviews with clear reporting on performance and incidents

- Defined change management aligned with ITIL best practices

Knowledge Transfer

We embed knowledge sharing throughout our engagement: - Documented runbooks and operational procedures - Regular knowledge sharing sessions with your team - Post-incident reviews that share learnings

Exit Process

When the engagement ends, we ensure an orderly transition: - Complete documentation of all configurations and procedures - Knowledge transfer sessions as required - 90-day standard notice period - No disruption to your services during handover

7 Why Appvia

We're a UK-based SME with deep expertise in cloud platforms and a strong track record delivering for UK government.

Cloud Partnerships

- AWS Advanced Tier Services Partner & Public Sector Partner
- Microsoft Azure Solutions Partner
- Google Cloud Partner

Open Source Credentials

- Creators of Terranetes (CNCF project for Terraform on Kubernetes)
- Developers of Wayfinder, a self-service cloud platform

Public Sector Experience

We understand the unique requirements of government and public sector organisations, from security clearance to GDS service standards.

"Appvia have helped us realise our cloud-first strategy with excellent product solutions. The security, governance and peace of mind brought in by the cloud landing zone has dramatically accelerated our journey to cloud-native adoption."

— **Peter Macchiocci**, Bank of England, Head of Platform Engineering

8 Contact Us

To discuss your requirements or request a detailed proposal, contact our Public Sector Team.

Phone: 020 3917 2888

Email: sales@appvia.io

Website: www.appvia.io

When you get in touch, please include: 1. The name of this service 2. Your organisation name 3. Your contact details 4. A brief description of your requirements

We'll arrange a discovery call to understand your needs and provide a tailored response.

Appvia helps organisations build, secure, and innovate across cloud, data platforms, and AI by blending practical engineering expertise with real-world experience.



About Appvia

Appvia helps organisations build, secure, and innovate across cloud, data platforms, and AI by blending practical engineering expertise with real-world experience. Drawing on our background in creating open source CNCF projects like Terranetes and our self-service developer platform, Wayfinder, we bring that industry know-how and expertise into your organisation to help unlock the true value of the Cloud. As an AWS Advanced Tier Partner and a Microsoft Solutions Partner, we specialise in delivering secure, scalable cloud-native platforms and services, including containerisation, FinOps, modernisation, and AI integration.

Build, secure, and innovate in the cloud | www.appvia.io



This document contains information that may be privileged or confidential and is the property of Appvia Ltd.

Public Copyright 2026 Appvia Ltd. All rights reserved.