



Crown
Commercial
Service

zensar



Zensar's Service Definition Document for AI-Driven Enterprise Knowledge Management and Intelligent Search

January 2026

Disclaimer & Confidentiality

The information in this document shall not be disclosed outside the Crown Commercial Service and shall not be duplicated, used or disclosed, in whole or in part, for any purpose other than to evaluate the document. This restriction does not limit the right of the Crown Commercial Service to use information contained in this document if it is obtained from another source without restriction. This proposal is valid up to 90 days from the date of submission of the proposal.

Table of Contents

1	Service Overview.....	1
2	Technical Architecture & Agentic Capabilities	1
3	Key Features	1
4	Onboarding and Offboarding	2
5	Key Notes.....	2

1 Service Overview

Zensar's AI-Driven Knowledge Management service transforms fragmented public sector data into a unified, intelligent resource. Aligned with the **Government AI Playbook**, the solution is designed to be **cloud-agnostic**, capable of deploying Microsoft Azure, AWS or Google Cloud. It addresses the "Digital Backbone" requirement by connecting disparate repositories (SharePoint, S3, Drive) into a single semantic search interface, utilizing **Agentic AI** to perform autonomous research and synthesis tasks.

2 Technical Architecture & Agentic Capabilities

Agentic AI Framework: Our solution moves beyond simple chatbots to deploy **autonomous agents** as defined in the Government AI Playbook. These agents perceive user intent, select appropriate tools (e.g., "Search Policy Database," "Summarise Report") and execute multi-step reasoning.

- **Implementation Flexibility:** Agents can be built using **Azure OpenAI Assistants**, **AWS Bedrock Agents** or **Google Vertex AI Agent Builder**, depending on the client's preference.

Integration with Current Investments (The "Digital Backbone"): We advocate a "layer, don't replace" strategy to maximize ROI on existing investments:

- **Unified Data Ingestion:** Automated pipelines ingest content from **Microsoft 365/SharePoint**, **AWS S3** or **Google Drive**. We utilise event-driven architecture (e.g., Azure Event Grid, EventBridge or Pub/Sub) to trigger processing when documents are updated.
- **Workflow Preservation:** The solution integrates existing governance tools (e.g., **Power Automate** or **AODocs**), ensuring AI enhances rather than bypasses established approval workflows.
- **Model Agnosticism:** The architecture supports a "Model Garden" approach, allowing departments to switch between models (e.g., GPT-4, Claude 3, Llama 3) based on performance, cost or sovereignty requirements.

Sovereignty & Security:

- **Sovereign Hosting:** For critical national functions, we support private model hosting and local RAG (Retrieval-Augmented Generation) implementations to ensure data never leaves the client's-controlled environment.
- **Identity Integration:** Security is managed via existing Identity Providers (e.g., **Microsoft Entra ID** or **Google IAM**), enforcing Zero Trust principles and granular access controls.

3 Key Features


- **Automated Taxonomy:** AI extracts document types and metadata automatically, populating enterprise data catalogs.
- **Deep Research Mode:** Agents autonomously synthesise answers from multiple documents, generating citations to ensure verifiability.
- **De-Duplication Engine:** SHA-256 hashing and semantic analysis identify duplicate content across the estate to reduce storage costs and "dark data".
- **Analytics & Reporting:** Integration with **Power BI**, **Looker** or **QuickSight** to track search metrics and content gaps.

4 Onboarding and Offboarding

- **Onboarding:** A structured Discovery phase maps the data estate and selects the optimal cloud landing zone. We configure "Kickstarter" pilots (e.g., Chat Interfaces) within 6-12 weeks.
- **Offboarding:** All metadata and vector indices are exported by open standards (JSON/CSV). Infrastructure is provisioned via Terraform/IaC to ensure portability and prevent vendor lock-in.

5 Key Notes

- **Definitions** "Service" refers to the AI-Driven Enterprise Knowledge Management solution, deployable on the Client's chosen cloud platform (Azure, AWS or GCP).
- **Intellectual Property Rights (IPR)** Zensar retains ownership of background IPR and methodologies. Custom configurations, specifically those integrating with the Client's specific "Digital Backbone" or proprietary data models, shall be assigned to the Client upon payment.
- **Liability** Zensar's liability shall be limited to 100% of the contract value or £1,000,000, whichever is lower. We accept no liability for the accuracy of outputs generated by third-party Foundation Models (hallucinations), but provide mitigation strategies (RAG, citation) as part of the service.
- **Data Protection and Sovereignty** We comply with the UK GDPR. Data processing locations are user-selectable to ensure compliance with **AI Sovereignty** requirements. We support deployment into UK-sovereign cloud regions to meet data residency obligations.
- **Service Level Agreement (SLA)** Service availability is targeted at 99.9%. Support levels (Standard, Gold, Platinum) define response times for P1-P4 incidents, ranging from 15 minutes to 24 hours.



We conceptualize, build and manage digital products through experience design, data engineering and advanced analytics for over 130 leading companies. Our solutions leverage industry-leading platforms to help our clients be competitive, agile and disruptive while moving with velocity through change and opportunity.

With headquarters in Pune, India, our 10,000+ associates work across 33 locations, including San Jose, Seattle, Princeton, Cape Town, London, Singapore and Mexico City.

For more information please contact: marketing@zensar.com | www.zensar.com

An  Company