



Zensar's Service Definition Document for AI-Driven Enterprise Strategy Evaluation and Simulation

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	Introduction.....	1
1.1	Service Scope.....	1
1.2	Experience Intelligence Architecture (Technical and Operational Overview)	1
1.2.1	Layer 1: Data Ingestion (Senses)	1
1.2.2	Layer 2: Core AI Engines (Brain).....	2
1.2.3	Layer 3: Platform Services (Tools).....	2
1.2.4	Layer 4: Human in the Loop Operations (Conscience)	2
1.2.5	Layer 5: Output and Action (Hands)	2
1.3	Alignment with the UK Government AI Playbook Principles.....	2
1.3.1	Know what AI is and its limitations.....	2
1.3.2	Use AI lawfully, ethically and responsibly	3
1.3.3	Know how to use AI securely.....	3
1.3.4	Meaningful human control	3
1.3.5	Manage the AI lifecycle	3
1.3.6	Use the right tool for the job.....	3
1.3.7	Open and collaborative.....	3
1.3.8	Work with commercial colleagues.....	3
1.3.9	Skills and expertise	3
1.3.10	Assurance and policies.....	3
1.4	Technical Sovereignty and Cloud Infrastructure	3
1.5	Implementation Methodology and Onboarding (Scan to Pilot to Scale)	4
1.6	Operational Model.....	4
1.7	Data Security and Privacy Assurance	4
1.8	Integrations and Interfaces	5
1.9	Roles and Responsibilities (Summary).....	5
1.10	Deliverables by Phase.....	5
1.11	Success Measures (KPIs and XLAs).....	5

1 Introduction

This service helps departments evaluate, de-risk and optimise enterprise strategy using AI-driven simulation and live experience intelligence. It combines Experience Twins (generative persona simulations) with Empathy Ops (continuous sensing) to forecast outcomes, detect strategic drift and trigger targeted micro-interventions. Human in the Loop (HITL) governance is applied at all critical stages.

Outcomes

- Evidence-based strategic decisions before broad rollout through a virtual pilot
- Real-time insight into citizen and employee sentiment and adoption risk
- Targeted nudges, risk maps and leadership dashboards for continuous improvement
- Cloud-agnostic, secure deployment within the buyer's sovereign tenant

1.1 Service Scope

In scope

- Experience intelligence architecture deployment across five layers
- Data connector configuration and secure integration with approved systems
- Experience Twin simulations, scenario experimentation and calibration
- Continuous sensing, sentiment and emotion analysis and narrative alignment detection
- HITL consulting, ethical oversight and decision support
- Nudge orchestration and organisational risk mapping
- Training and enablement for Twin Designer and operational runbooks
- Ongoing model health checks and lifecycle management

Out of scope (available through change control)

- Integrations with unapproved data sources
- Non-standard custom model builds outside platform patterns
- Broad organisational change programmes beyond advisory and nudges
- Retrospective data cleansing or data quality remediation projects

1.2 Experience Intelligence Architecture (Technical and Operational Overview)

The Experience Twin and Empathy Ops platform is a modular microservices architecture designed for flexibility, security and scale. It comprises five logical layers.

1.2.1 Layer 1: Data Ingestion (Senses)

Function

Secure, cloud-agnostic collection of structured and unstructured signals through Experience Data Connectors.

Supported sources

- Interaction data such as anonymised call transcripts, chat logs and email threads

- Feedback data such as CSAT and employee pulse surveys including Glint and Culture Amp
- Operational data such as digital telemetry including clickstreams and dwell time
- Contextual data such as policy documents, intranet pages and strategy PDFs

Design

Standard RESTful APIs compatible with major cloud providers and commonly used SaaS in government.

1.2.2 Layer 2: Core AI Engines (Brain)

- **Persona Simulation Engine** builds digital twins of user cohorts with generative AI and uses retrieval-augmented generation to ground simulations in departmental knowledge and reduce hallucinations
- **Sentiment and Emotion Engine** use specialised NLP to detect nuanced emotional states in government contexts such as resignation and anger
- **Narrative Alignment Engine** measures semantic distance between strategy and operations to surface strategic drift

1.2.3 Layer 3: Platform Services (Tools)

- **Experience Twin Designer Studio** offers low-code configuration of simulations with variable controls such as age and digital literacy for policy stress testing
- **Scenario Sandbox** provides a secure and isolated environment for large-scale simulations without impact on live systems

1.2.4 Layer 4: Human in the Loop Operations (Conscience)

- **Consultant Interpretation** by expert strategists to derive second-order insights such as cultural and technical resistance
- **Ethical Oversight** with workflows to assess simulations for bias before leadership consumption

1.2.5 Layer 5: Output and Action (Hands)

- **Nudge Engine** for targeted micro-interventions by email, Slack or in-app prompts
- **Risk Maps** as visual heatmaps to highlight emotional friction and adoption risk across the organisation

1.3 Alignment with the UK Government AI Playbook Principles

1.3.1 Know what AI is and its limitations

Simulations expose confidence intervals and position twins as probabilistic advisors to reduce automation bias.

1.3.2 Use AI lawfully, ethically and responsibly

Fairness filters test for bias related to protected characteristics and GDPR data minimisation is applied with synthetic data where feasible.

1.3.3 Know how to use AI securely

Secure by design with bring your own key encryption. Data remains in the buyer's sovereign cloud tenant.

1.3.4 Meaningful human control

The HITL layer prevents significant decisions without human review and approval.

1.3.5 Manage the AI lifecycle

Model drift monitoring and retraining handle changes such as post-election behaviour.

1.3.6 Use the right tool for the job

AI is used for behavioural simulation and pattern detection. Deterministic tasks use standard logic for efficiency.

1.3.7 Open and collaborative

Open APIs using OpenAPI 3.0 enable cross-department sharing including with CDDO.

1.3.8 Work with commercial colleagues

Dashboards link experience improvements to cost savings such as reduced handling time.

1.3.9 Skills and expertise

Lot 3 training and Twin Designer certification are available for civil servants.

1.3.10 Assurance and policies

The service supports the Algorithmic Transparency Recording Standard for accountable explanations.

1.4 Technical Sovereignty and Cloud Infrastructure

All components are packaged as Docker containers for portability and consistency. Deployment uses Kubernetes and supports EKS, AKS, GKE and OpenShift for on-premise needs. The platform uses multiple data backends such as PostgreSQL and MySQL and vector stores such as Pinecone and Milvus so it can use the buyer's managed services. This enables migration across cloud providers with minimal refactoring and helps avoid vendor lock-in.

1.5 Implementation Methodology and Onboarding (Scan to Pilot to Scale)

A five-phase approach delivers early value and manages risk.

Phase 1: Discovery and Design (Weeks 1 to 3)

Activities listening workshops, data landscape mapping and persona definition

Outputs listening architecture blueprint and persona definition document

Compliance initial DPIA and AI ethics assessment

Phase 2: Prototype and MVP Configuration (Weeks 4 to 8)

Activities deploy the containerised platform in the buyer cloud for dev and test, configure ingestion pipelines and train initial twins

Output working MVP with active ingestion and basic simulations

Phase 3: Pilot and Calibration (Weeks 9 to 12)

Activities run in shadow mode, compare predictions to actuals and calibrate

Output calibration report that validates accuracy

Phase 4: Production and Training (Weeks 13 to 18)

Activities promote to production, roll out the Empathy Ops dashboard and deliver Twin Designer training

Output live integrated system and a trained internal team

Phase 5: Commercial Launch and Managed Service (Ongoing)

Activities transition to BAU, ongoing insight advisory and quarterly model health checks for drift

Output continuous strategic value and experience intelligence

1.6 Operational Model

- Operating rhythm with weekly insight reviews, monthly value realisation forums and quarterly strategy refresh
- HITL gates with ethics review before mass communications and leadership sign-off for policy-level changes
- Model lifecycle with versioned models, drift monitoring, retraining windows and rollback plans
- NudgeOps with a library of micro-interventions, audience targeting, opt-out controls and impact measurement
- Risk management with a pre-production scenario sandbox, staged feature flags and progressive rollout

1.7 Data Security and Privacy Assurance

Operations run within the buyer's sovereign cloud tenant. No production data leaves the tenant. Data is secured in transit using TLS 1.2 or higher and at rest using bring your own key. Role-based access control with least privilege and documented break-glass procedures is enforced. PII minimisation and pseudonymisation are applied for analytics with retention schedules aligned to

buyer policy. The service maintains event logs, model inference logs and ethical assessment records to support full transparency.

1.8 Integrations and Interfaces

- **Inbound** REST connectors for interaction, feedback, operational and contextual data
- **Outbound** OpenAPI 3.0 services for dashboards, risk maps, nudges and data exchange with enterprise tools
- **Toolchain examples** voice of the customer platforms, ITSM, collaboration suites and data lakes enabled through standard APIs
- **Non-functional** backpressure controls, idempotency and fault-tolerant ingestion pipelines

1.9 Roles and Responsibilities (Summary)

- **Buyer Executive Sponsor** strategy ownership, funding and governance
- **Buyer Product Owner** backlog, prioritisation, UAT and acceptance
- **Zensar Engagement Lead** delivery oversight, risk and commercial governance
- **AI Ethics Lead** fairness filters, DPIA, bias audits and assurance documentation
- **Solution or Platform Architect** secure by design cloud-agnostic deployment
- **Data and ML Engineers** connectors, pipelines, model training and monitoring
- **Twin Designers** scenario design, parameterisation and experiment orchestration
- **Insight Consultants** interpretation, recommendations and nudge playbooks

A detailed RACI is provided during onboarding.


1.10 Deliverables by Phase

- **Discovery** listening architecture blueprint, persona definitions, DPIA and initial ethics assessments
- **MVP** configured platform, connectors, baseline twin models and runbooks
- **Pilot** calibration report, model validation metrics and drift baselines
- **Production** Empathy Ops dashboard, nudge library, operating playbooks and training completion
- **BAU** quarterly model health checks, strategy insight packs and transparency records

1.11 Success Measures (KPIs and XLAs)

- Calibration accuracy for the delta between predicted and actual sentiment or adoption
- Time to insight from signal ingestion to executive-ready recommendations
- Drift detection time and retraining lead time
- Nudge impact such as reduction in friction hot spots and uplift in adoption
- Narrative alignment score trends
- Ethical compliance including bias audit pass rate and completeness of transparency artefacts

Thresholds and service levels are baselined during transition and aligned to buyer standards.



We conceptualise, 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