



Zensar Service Definition Document for Data Analytics & AI

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	Data Analytics & Artificial Intelligence Services Overview	1
2	Service(s) Overview	1
2.1	Azure Data Platform Implementation and Support.....	1
2.2	Migrate legacy DWH/ETL to Snowflake, Databricks	2
2.3	Service: Implement Databricks Lakehouse for analytics and ML	3
2.4	Service: Snowflake Data Warehouse Build, Migration and Optimisation.....	4
2.5	Managed Support for Snowflake/Databricks/Azure Data Platforms	5
2.6	Gen AI based data mapping, extraction, assessment & validation.....	6
2.7	Annotation as a service with Scalable human review and AI	6
2.8	Agentic AI based data exploration, insight generation	7

1 Data Analytics & Artificial Intelligence Services Overview

Zensar's team of over 1200+ Data Engineering experts, with more than 80% certified across leading technologies like AWS, Databricks, Snowflake, Azure and GCP, stands as a beacon for businesses seeking unparalleled data solutions.

We offer services such as designing robust data strategies, recommending cloud services & migrating data platforms to setting up data foundations and constructing intricate pipelines with complex transformations, creating reports to generate actionable intelligence, automating processes to increase productivity and AI-GenAI solutions for predictive analytics.

Zensar has been a catalyst in numerous data modernisation journeys. With its technical proficiency, execution experience and army of accelerators and frameworks for data ingestion, transformation, pre-built data marts and KPIs, Zensar ensures innovation at every step, propelling businesses toward data excellence. We have invested & built ZenseAI.Data, a platform that brings generative and agentic AI straight into the data lifecycle, so companies don't have to do every task manually anymore - from ingesting and cleaning data to building analytics and transformation pipelines.

With AI accelerating everything from code generation to quality checks, ZenseAI.Data lets team handle data modernisation at scale while reducing time and cost. Under the hood, the platform is powered by gen AI and intelligent agents that can parse legacy SQL and ETL logic, regenerate optimised assets and help automate complex engineering tasks. It also supports governance, metadata enrichment and analytics workflows that make data easier to explore and use across teams. This tech foundation reflects the shift in enterprise needs - from siloed engineering work to automated, insight-driven pipelines that scale with business demand. In essence, ZenseAI.Data is Zensar's response to the growing demand for **AI-assisted data modernisation**, faster insights, automated engineering and a smoother bridge from old systems to modern cloud data architectures.

2 Service(s) Overview

2.1 Azure Data Platform Implementation and Support

We design, build and support data platforms using Microsoft Fabric and Azure Data Lake Storage (ADLS). We implement secure platform foundations, ingest data into the lake, build curated models and deliver governed BI through Fabric (including semantic models and reports where required). This service provides the activities required to implement or improve a Fabric-based analytics platform with ADLS as the underlying data lake. It covers platform setup, ingestion patterns, transformation and modelling and operational readiness. The service can be used for new platform delivery, migration from legacy BI/warehouse solutions, or stabilisation and optimisation of an existing Fabric estate.

Key Features:

- Discovery workshops (use-cases, data domains, SLAs, security and operating model)
- Azure platform foundations for ADLS (subscriptions, networking, identity and access)
- ADLS setup (storage accounts, containers, access patterns, lifecycle policies)

- Microsoft Fabric tenant/workspace setup and configuration (capacity considerations, workspace structure)
- Data ingestion into the lake and Fabric (batch and scheduled ingestion; patterns agreed per source)
- Transformation and curation using Fabric capabilities (medallion/curated layer patterns)
- Data modelling for analytics (semantic modelling approach and naming standards)
- BI enablement in Fabric (datasets/semantic models; report delivery where agreed)
- Data quality checks and reconciliation (rules, exception handling, monitoring)
- Security controls (RBAC, least privilege, audit logging; customer policies applied)
- Monitoring and operational readiness (runbooks, support processes, handover)
- Cost and capacity optimisation guidance (Fabric capacity utilisation, storage optimisation)

Benefits:

- Faster delivery of a governed Fabric analytics platform
- ADLS-based foundation supports scalable and cost-effective storage.
- Consistent curated datasets and semantic models for reliable reporting
- Improved operational control through monitoring and documented runbooks.
- Clear approach to Fabric capacity management and optimisation

2.2 Migrate legacy DWH/ETL to Snowflake, Databricks

We help organisations migrate legacy data warehouses and ETL workloads to a modern cloud data platform using Snowflake and/or Databricks. We assess current data estate, design the target architecture, migrate data and pipelines, validate results and support cutover to production. This service supports the discovery, planning, build, migration and transition activities needed to move from on-premises or legacy DWH/ETL tools to Snowflake and Databricks. We have proprietary accelerator – ZenseAI.data which is Zensar’s response to the growing demand for AI-assisted data modernisation, faster insights, automated engineering and a smoother bridge from old systems to modern cloud data architectures. It solves the very real pain points of slow delivery, fragmented data and manual processes that have held enterprises back from fully realising AI’s potential. We focus on secure delivery, data quality and repeatable deployment and operations through Zensar proprietary accelerators.

Key Features:

- Legacy platform assessment (data sources, ETL jobs, schedules, SLAs, dependencies)
- Migration approach and plan (phased, domain-by-domain, or workload-based)
- Target architecture design for Snowflake and/or Databricks on any cloud
- Data ingestion design and build (batch, near-real-time)
- ETL/ELT migration and refactoring (including performance optimisation)
- Data model migration support (dimensional, data vault, or curated layer patterns)
- Data reconciliation and validation (row counts, totals, sampling, data quality checks)
- Cutover planning and execution (parallel run, rollback approach, go-live support)

- Operational readiness (monitoring, alerting, runbooks, handover)
- Knowledge transfer workshops for customer teams

Benefits:

- Reduced legacy platform risk and maintenance overhead.
- Improved scalability and performance for analytics workloads
- Clear migration plan with measurable checkpoints and governance
- Better data quality controls and repeatable deployments
- Enables modern analytics and ML use-cases on curated data.

Skills – Azure, MS Fabric, 150+ certifications

2.3 Service: Implement Databricks Lakehouse for analytics and ML

We design and implement Databricks Lakehouse on Azure to support analytics and machine learning. We set up the environment, establish secure governance, build ingestion and transformation pipelines and deliver curated datasets and ML-ready features aligned to your use-cases. This service delivers the technical build and enablement activities required to implement a governed Databricks platform and a Lakehouse data architecture. It supports both new platforms and improvements to existing Databricks estates. We have proprietary accelerator – ZenseAI.data which is Zensar's response to the growing demand for AI-assisted data modernisation, faster insights, automated engineering and a smoother bridge from old systems to modern cloud data architectures. It solves the very real pain points of slow delivery, fragmented data and manual processes that have held enterprises back from fully realising AI's potential.

Key Features:

- Use-case and requirements workshops (analytics, reporting, ML, operational needs)
- Azure and Databricks architecture design (networking, identity, access controls)
- Workspace provisioning and configuration
- Data governance setup using unity catalog (cataloguing and access patterns, where required)
- Data ingestion pipelines from source systems to bronze/silver/gold layers
- Data transformation and optimisation for analytics workloads
- Feature engineering support for ML use-cases (where required)
- Job orchestration, monitoring and alerting
- CI/CD enablement for notebooks/jobs
- Performance tuning and cost optimisation guidance
- Operational readiness, runbooks and handover/training

Benefits:

- Faster delivery of analytics-ready curated datasets

- Secure, governed access to data for multiple teams
- Reduced time to deploy and operate Databricks workloads.
- Improved reliability through monitoring and operational processes
- Platform built to support ML development and deployment workflows.

Skills – Databricks Data Engineer, Databricks unity catalog SME, Databricks Mosaic AI

Zensar has 200+ people in Databricks out of which 150+ are Databricks certified.

2.4 Service: Snowflake Data Warehouse Build, Migration and Optimisation

We design, build and migrate data warehouses on Snowflake and optimise existing Snowflake environments for performance, reliability and cost. We support platform setup, ingestion and transformation patterns, data modelling, security controls, testing and cutover. This service delivers the engineering activities required to implement Snowflake as a cloud data warehouse, migrate workloads from legacy platforms, or improve an existing Snowflake estate. We have proprietary accelerator – ZenseAI.data which is Zensar's response to the growing demand for AI-assisted data modernisation, faster insights, automated engineering and a smoother bridge from old systems to modern cloud data architectures. It solves the very real pain points of slow delivery, fragmented data and manual processes that have held enterprises back from fully realising AI's potential. We work with your teams to define target architecture, implement secure access controls, migrate data and pipelines and establish operational processes for a production service.

Key Features:

- Discovery and assessment (current-state DWH/ETL, workloads, SLAs, dependencies)
- Target architecture and migration approach (phased migration, parallel run, rollback planning)
- Snowflake account/environment setup guidance (roles, warehouses, databases, network controls)
- Data ingestion design and build (batch and scheduled loads; patterns agreed per source)
- ELT implementation (SQL-based transformations; modelling approach agreed)
- Data modelling support (dimensional/curated layer patterns, documentation standards)
- Data quality checks, reconciliation and testing (automated checks where agreed)
- Performance tuning (warehouse sizing, clustering, query optimisation, caching strategy)
- Cost optimisation (warehouse governance, schedules, resource monitors, usage reporting)
- Security controls (RBAC, least privilege, auditing; customer policies applied)
- Cutover execution and go-live support (HyperCar period agreed)
- Operational readiness (monitoring, runbooks, incident/change processes, handover)

Benefits:

- Faster and lower-risk migration to a modern cloud data warehouse
- Improved reporting performance and workload stability
- Better cost control through Snowflake usage governance

- Clear security model and audit-friendly operational processes
- Repeatable delivery approach with documented artefacts

Skills – Snowflake data engineer, Snowflake cortex SME

We have 100+ team in Snowflake technology with 60+ certified.

2.5 Managed Support for Snowflake/Databricks/Azure Data Platforms

We provide ongoing managed support for Snowflake, Databricks and Azure-based data platforms. This includes incident response, monitoring and operational improvements, minor enhancements, performance and cost optimisation and platform health management aligned to agreed service levels. This service supports customers operating production data platforms and data pipelines. We work within an agreed support model to maintain platform stability, resolve incidents, manage operational tasks and implement continuous improvements. We also have one of our accelerators – ZenseAI.data – AMS which helps reduce cost in application management services by leveraging AI capabilities in AMS. The service can cover one or more platforms (Snowflake, Databricks, Azure services) and can be tailored to your operating model.

Key Features:

- Service onboarding and transition (service inventory, access setup, support runbooks)
- Incident management (triage, investigation, mitigation and post-incident reviews)
- Request fulfilment (access requests, configuration changes, standard operational tasks)
- Monitoring and alerting setup/improvement (platform and pipeline monitoring)
- Ticket similarity engine, Smart Ticket Assignment & Ticket Volume Forecast
- SLA Breach Propensity, Team Utilisation Insights & User FAQ Bot
- Platform health checks (reliability, security posture, performance and cost)
- Performance optimisation (query/job tuning, capacity/cluster/warehouse optimisation)
- Cost optimisation (usage governance, scheduling, tagging/budgets where applicable)
- Release and change support (deployment coordination, rollback planning, release notes)
- Data pipeline support (failure handling, reprocessing, SLA monitoring)
- Security support (role reviews, least privilege, audit log reviews, policy alignment)
- Documentation maintenance (runbooks, known error database, operational procedures)
- Monthly service reporting (KPIs such as incident volume, SLA attainment, cost trends)

Benefits:

- Improved reliability and reduced downtime for critical data services
- Faster incident resolution through standard triage and runbooks
- Ongoing cost control and performance improvements
- Clear accountability via agreed support processes and reporting
- Frees customer teams to focus on delivery rather than operations.

2.6 Gen AI based data mapping, extraction, assessment & validation

We use generative AI, supported by engineering controls and human review where required, to accelerate data mapping, extraction, assessment and validation activities. This service helps organisations understand data assets, document mappings between systems, extract structured fields from documents and validate migrated or integrated datasets. With AI accelerating everything from code generation to quality checks, we have built an accelerator, ZenseAI.Data which lets team handle data modernisation at scale while reducing time and cost. Under the hood, the platform is powered by gen AI and intelligent agents that can parse legacy SQL and ETL logic, regenerate optimised assets and help automate complex engineering tasks. It also supports governance, metadata enrichment and analytics workflows that make data easier to explore and use across teams. We apply GenAI techniques to speed up repetitive analysis and documentation tasks while maintaining auditability through defined workflows, testing and quality gates. Outputs are provided as structured artefacts that can be reviewed and approved by customer data owners.

Key Features:

- Data discovery workshops (scope, data domains, success criteria, controls)
- Automated profiling and assessment of source data (completeness, uniqueness, anomalies)
- GenAI-assisted mapping documentation (source → target field mapping drafts)
- Rules extraction support (e.g., transformation logic and business rules, where available)
- Tech moderniser - Migrating from SAS/SSIS to Cloud
- GenAI-assisted document extraction (structured fields from semi-structured documents)
- Validation framework setup (reconciliation checks, sampling approach, exception handling)
- Migration validation support (row counts, totals, aggregates, key integrity checks)
- Human-in-the-loop review workflow (approval steps, escalation, sign-off)
- Prompting and evaluation approach (test sets, acceptance thresholds, drift checks)
- Security and access controls aligned to customer policies (data minimisation where possible)
- Documentation pack suitable for governance/audit (method, outputs, decisions)

Benefits:

- Faster production of mapping and assessment artefacts
- Reduced manual effort in extraction and validation tasks.
- Improved consistency via repeatable validation rules and templates
- Clear traceability of mappings, exceptions and approvals
- Supports higher confidence in migration and integration outcomes.

2.7 Annotation as a service with Scalable human review and AI

We provide managed data annotation services combining a scalable human workforce with automation and quality controls. We label and review datasets to support machine learning and analytics, including text, images, audio and document data, delivering outputs in agreed schemas with measurable quality. This service provides an end-to-end annotation operating model: onboarding, guideline creation, annotation execution, quality assurance and delivery. Automation is used to improve throughput with human reviewers providing final quality gates.

Key Features:

- Multi-Modal Data Annotation
- Text (NER, sentiment, classification, Q&A)
- Image (bounding boxes, segmentation, object detection)
- Audio (transcription, intent recognition)
- Video (frame-by-frame tagging)
- Tabular data (classification, normalisation)
- Domain-Specific Labelling
- High-Tech ILM, legal contracts, healthcare, retail, automotive
- Pre-trained taxonomies and ontologies to accelerate annotation.
- Annotation Platform & Tooling
- Custom annotation workflows using open-source or proprietary tools Plugin support for human-in-the-loop feedback.
- Model Feedback Loop Integration
- Model performance monitoring to inform active learning and data prioritisation.
- Annotation Ops (AnnotationOps) as a Service
- Scalable human workforce + automation for enterprise AI/ML teams
- Managed services or flexible staff augmentation

Benefits:

- Scales annotation capacity to meet delivery deadlines.
- Improves model training quality through controlled QA and adjudication.
- Reduces total cost of labelling with automation-assisted workflows.
- Provides predictable delivery through measured throughput and reporting.
- Enables repeatable, auditable annotation processes suitable for governed environments.

2.8 Agentic AI based data exploration, insight generation.

We implement agentic AI workflows to help users explore data, generate insights and produce analysis outputs faster, using governed access to your enterprise data. The service combines AI agents with agreed guardrails, human oversight and auditable workflows to support analytics teams and business users. The service is designed to operate within customer security policies, with controlled data access, logging and validation of outputs.


Key Features:

- Discovery workshops (use-cases, user groups, data domains, risk controls, success measures)
- Selection of operating model (self-service assistant, analyst copilot, or governed request workflow)
- Integration with enterprise data sources (curated warehouse/Lakehouse datasets)

- Semantic layer and metrics definition support (to reduce ambiguity and improve answer quality)
- Knowledge Explorer
- Auto data visualiser & Narrative Insights
- Report Migrator Agent & Report Usage Agent
- Agent workflow design (tool use, step-by-step reasoning, validation steps, escalation paths)
- Guardrails and controls (role-based access, restricted actions, data minimisation, policy rules)
- Retrieval configuration (approved datasets, documentation and data dictionary integration)
- Output validation patterns (cross-checks, confidence thresholds, human review for key outputs)
- Auditability (logging of prompts, actions, queries executed and outputs generated)
- Pilot implementation and user testing (acceptance criteria, feedback cycles, refinement)

Benefits:

- Faster insight generation for analysts and decision makers
- Improved consistency through governed datasets and defined metrics
- Reduced manual effort for routine exploration and reporting tasks.
- Controlled and auditable use of AI against organisational data
- Supports adoption while managing security and quality risks.



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