MDDI - Metadata Driven Data Integration

SERVICE DEFINITION

U ▪
S T

ust.com

Service Overview

MDDI (Metadata-Driven Data Integration) is a reusable framework that acquires, validates, standardizes, and stores data across on-prem and cloud systems with full automation and telemetry. It orchestrates pipelines using ADF, Databricks/Spark, and PowerShell triggers, enabling rapid onboarding of sources, consistent quality, and governed, audit-ready data delivery to Data Lakes, Lakehouses, Delta, and Data Warehouses.

Service Scope – Platform Capabilities

- Connectivity & Ingestion: Databases, files, on-prem/cloud storage; batch, event-based, and scheduled ingestion.
- Pipelines: Metadata-driven recipes (table/JSON).
- Transform & Standardize: Schema alignment, enrichment, SCD/CDC patterns, and reusable transformation libraries.
- Targets: Data Lake, Lakehouse, Delta Lake, and Data Warehouse with partitioning/compaction strategies.
- Orchestration: End-to-end orchestration with retries, dependencies, and SLA-aware execution.
- Telemetry: End-to-end logging, metrics, lineage hooks, and audit trails.
- Security: RBAC/least-privilege, key/secrets management, encryption in transit/at rest.
- Consumption: BI/analytics, APIs, and direct SQL endpoints with curated zones.

Service Scope – Operational Services

- Environment Setup: Provisioning, networking, identity (SSO/MFA), secrets, key management.
- Pipeline Operations: Scheduling, monitoring, alerts, failure handling, and SLA governance.
- Performance Management: Spark tuning, autoscaling, file layout optimization (partitioning/z-order).
- Data Quality & Observability: Validations, profiling, schema drift detection, quarantine/replay.
- Release & Change: CI/CD, versioning of recipes/notebooks, controlled change windows.
- Backup & Recovery: Configurable retention, restore testing, disaster-recovery runbooks.
- Governance Operations: Access reviews, audit reports, policy enforcement hooks.
- Support: Incident/request/problem management with defined SLAs.

Out of Scope

- Custom application development beyond MDDI patterns.
- Third-party licenses, cloud infrastructure fees, and data egress charges.
- Source-system remediation or upstream data cleansing.
- On-prem hardware procurement/maintenance.
- Bespoke ML/model development not tied to ingestion use cases.
- Extended hypercare, on-site, or 24×7 coverage unless contracted.

Customer Responsibilities

- Provide access to sources, SMEs, network endpoints, and test data.
- Manage identity provider, user lifecycle, and role approvals.
- Own data classification, retention, and compliance guidance.
- Fund/manage cloud subscriptions and connectivity (VPN/private endpoints).
- Participate in UAT, sign-offs, and change approvals.
- Define acceptance criteria and success metrics per use case.

Service Process

- Discovery & Design: Scope sources/targets, mappings, security, and governance.
- Provisioning: Environments, connectivity, identity integration, secrets.
- Build: Recipes, ADF pipelines, Databricks/Spark notebooks; tests and SLAs.
- Validate: Functional/UAT, performance tuning, operational readiness.
- Go-Live: Controlled release, monitoring, initial hypercare.
- Operate: Continuous improvements, cost/performance optimization, governance reviews.
- Offboarding (if applicable): Data extraction, documentation handover, access revocation.

Support Hours

- Standard: Monday–Friday, 9:00–18:00 local customer time (excluding public holidays).
- Extended/24×7: Available as an add-on for production-critical workloads.

Support Process

- Intake: Authenticated portal/ticket/email; auto-acknowledgment with ticket ID.
- Triage: Severity assignment, routing, initial diagnostics.
- Resolution: Workaround/fix via CI/CD and change control.
- Communication: Status updates per SLA; incident reports for Sev-1/Sev-2.
- Closure: Customer confirmation, knowledge-base update, root-cause review (as applicable).

Defect Severity Levels

- Severity 1: Service unavailable, critical compliance impact
- Severity 2: Major functionality impaired, workaround available
- Severity 3: Minor issues, no regulatory impact
- Severity 4: Cosmetic or enhancement requests (SLAs customizable per contract/region.)

Training & Enablement

- Onboarding Workshops: Framework overview, governance, best practices.
- Role-Based Training: Data engineer (recipes/pipelines), platform admin (security/ops), analyst (consumption).
- Assets: Playbooks, quick-start guides, templates, sample recipes/pipelines.
- Office Hours: Scheduled Q&A; optional shadowing for first releases.
- Certification Path: Optional proficiency badges and hands-on labs.