SynkRhino AI

SERVICE DEFINITION

U ·
S T

ust.com

Service Overview

SynkRhino AI is an AI-powered data pipeline platform that automates profiling, modeling, PySpark code generation, orchestration, execution analytics, and data-quality validation through a multi-agent architecture (Profiling, Transformation, Analytics, DataOps, DQ Testing). It can be deployed on the customer's existing cloud subscription (AWS/Azure/GCP) or on services procured by SynkRhino per customer requirements.

Service Scope – Platform Capabilities

- Data Ingestion & Connectivity: File, object storage, databases, lakes/warehouses; schema inference and source authentication.
- Auto Data Profiling: Statistics, anomalies, patterns, completeness, uniqueness, cardinality, PII signals.
- AI Code Generation (PySpark): Natural-language requirements → optimized, standards-compliant PySpark with DQ checks, logging, and monitoring hooks.
- Orchestration & Scheduling: Prefect-based flows, retries, concurrency limits, task timeouts, SLAs, and alerting.
- Transformation & Modeling: Template-driven mappings, business rules, SCDs, joins, aggregations, partitioning, and performance tuning.
- Data Quality & Testing: Rule catalog, threshold policies, reconciliation, drift checks, quarantine/healing workflows.
- Execution Analytics: Run lineage, runtime metrics, success/failure rates, resource usage, KPI dashboards.
- APIs & Extensibility: REST APIs/SDKs, webhooks, plug-ins for custom connectors/transformers/validators.
- Security & Governance: RBAC, SSO/MFA, secrets vault bindings, audit logs, policy controls.
- Multi-cloud Ready: Region/AZ placement, cross-region replication options aligned to customer policy.

Service Scope – Operational Services

- Environment Setup: Landing zones, networking, identity, secrets, baseline guardrails.
- Onboarding & Configuration: Source integrations, rule catalog seeding, agent setup, initial pipelines.
- Run Operations: Monitoring, alert routing, incident triage, execution restarts, and routine housekeeping.
- Change Management: Versioned configuration-as-code, promotion (dev→test→prod), blue/green or canary releases.
- Observability: Dashboards, synthetic probes, log/metric/trace correlation, monthly operations report.
- Backup & DR (platform config): Scheduled backups of configs/metadata, restore procedures; customer data DR follows customer cloud policy.
- Security Ops: Access reviews, key rotation schedules for platform secrets, vulnerability patch windows.

Out of Scope

- Building non-platform custom applications, UIs, or unrelated integrations.
- Customer network, endpoint, or on-prem hardware administration.
- Third-party cloud/AI provider billing, quota management, or contract negotiation.
- Model training/fine-tuning services (unless contracted).
- Data stewardship, master data management, or business approval of results.
- Performance tuning of non-platform systems (DBs, DWHs) beyond connection best practices.

Customer Responsibilities

- Provide/maintain cloud subscriptions, credentials, network routes, and required quotas.

Confidential and Proprietary. © 2026 UST Global Inc

- Ensure data access permissions, governance policies, and compliance requirements are defined.
- Curate business rules, acceptance criteria, and data-quality thresholds.
- Operate non-platform components (sources/targets, firewalls, VPC/VNet, private links).
- Participate in change approvals, UAT, and production cutovers.
- Manage costs for cloud/AI usage under customer-owned subscriptions.
- Report incidents through agreed channels and maintain on-call contacts for escalations.

Service Process

- Initiation: Scope & success criteria; environment readiness checklist.
- Design & Enablement: Source mapping, rule catalog, orchestration and SLAs, access control.
- Implement: Configure agents, generate PySpark, build DQ suites, provision alerts.
- Test & Harden: Unit/integration/DQ tests; performance and DR drills.
- Go-Live: Controlled release, runbook activation, KPI baselines.
- Operate & Improve: Continuous monitoring, error-budget review, backlog grooming, and quarterly optimization.

Support Hours

- Standard: Monday–Friday, 09:00–18:00 (customer local business hours), excluding local holidays.
- Extended/Premium (optional): 24×7 for P1 incidents and scheduled releases (separately priced).

Support Process

- Channels: Ticketing portal/email; optional Slack/Teams bridge for P1/P2; status dashboard & API.
- Intake & Triage: Auto-acknowledge, severity assessment, assignment to on-call engineer.
- Comms Cadence: Per severity (see below), with ETA/next steps; RCA within 5 business days for P1 major incidents.
- Closure: Fix validation, knowledge-base entry, and preventive action tracking.

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

Note: Feature requests/enhancements are not covered by support SLAs and follow the change process.

Training & Enablement

- Foundational Enablement (included): Admin/Operator overview, pipeline configuration, DQ rules, troubleshooting (remote, 4–6 hours).
- Role-based Sessions (included): Data Engineer, Data Steward, Platform Admin quick-starts with labs and runbooks.
- Artifacts (included): Playbooks, configuration-as-code templates, best-practice checklists, and sample rule catalogs.
- Advanced Training (optional): Performance tuning, custom plug-in development, CI/CD, DR/Chaos drills, and enterprise governance workshops.

Confidential and Proprietary. © 2026 UST Global Inc