

FEATURE FACTORY — PRODUCT FACTSHEET

Machine learning for the masses.

Feature engineering, model training, scoring, deployment, and governance, all running inside your existing data warehouse. No data movement, no external ML platform, no specialist ML team required.

8

Engines: ML, governance, compliance, and more

0

Data movement required

3

Compliance frameworks built in: GDPR, CCPA, EU AI Act

Free

Community tier, source available

THE APPROACH

ML first. In-database. On the data you already have.

Most enterprise ML deployments start with a problem that has nothing to do with the model: getting data out of the warehouse and into an external ML platform. That movement is slow, expensive, and creates a separate governance problem.

Feature Factory takes a different approach. It runs feature engineering, model training, scoring, and deployment directly inside your existing database. Your data never moves, your governance controls remain in place, and your ML pipeline runs at warehouse scale without a separate platform.

EIGHT ENGINES

What Feature Factory provides

▪ Feature Engineering Engine

Define and compute ML features using in-database SQL pushdown, with full versioning and a feature store.

▪ Data Quality Engine

Automated table and column profiling with quality scoring across completeness, distribution, and pattern validation.

▪ Drift Detection Engine

Monitors feature and data drift using Population Stability Index metrics, with configurable alert thresholds.

▪ Geospatial Analytics Engine

Spatial feature engineering with distance calculations, region lookups, and UK postcode geocoding.

▪ Model Management & AutoML

Train, version, deploy, and run predictions in-database. AutoML recommends algorithms with full experiment tracking.

▪ Governance Engine

Business domain hierarchy, data stewardship, glossary, and lineage tracking, with impact analysis and a full audit trail.

▪ Compliance Automation Engine

GDPR, CCPA, and EU AI Act compliance built in: DSAR management, consent tracking, and automated erasure workflows.

▪ Demographic Engine

Enriches data with UK demographic profiles (population, age, income, OAC classification) and multi-country lookups.

Built for regulated industries. Compliance functionality is built into the platform from the ground up: DSAR management, consent tracking, Right to Be Forgotten erasure, ROPA and DPIA documentation, and EU AI Act system documentation, with a full tamper-evident audit trail.

EDITIONS & PRICING

Editions defined by database platform

Capability	Community	Professional	Enterprise
Database platform	PostgreSQL + MADlib	Cloudberry / Greenplum / Databricks	IBM Netezza + INZA
Availability	Roadmap	Roadmap	Shipping today
ML algorithm suite	Baseline (MADlib)	Full MADlib / Spark MLlib	Full INZA suite
AutoML	✗	✓	✓
Compliance automation (GDPR / CCPA / EU AI Act)	✗	✗	✓
Activation destinations (Salesforce, HubSpot, etc.)	✗	✗	✓
From (per month)	Free	£1,500	£5,000

Feature Factory ships today on IBM Netezza with INZA. The Community (PostgreSQL) and Professional (Cloudberry/Greenplum/Databricks) tiers are on the roadmap.

Bring ML to Your Data



Feature Factory ships today on IBM Netezza. Join the waitlist at smart-associates.biz/products/feature-factory to be notified as new platforms become available.

smart-associates.biz · info@smart-associates.biz