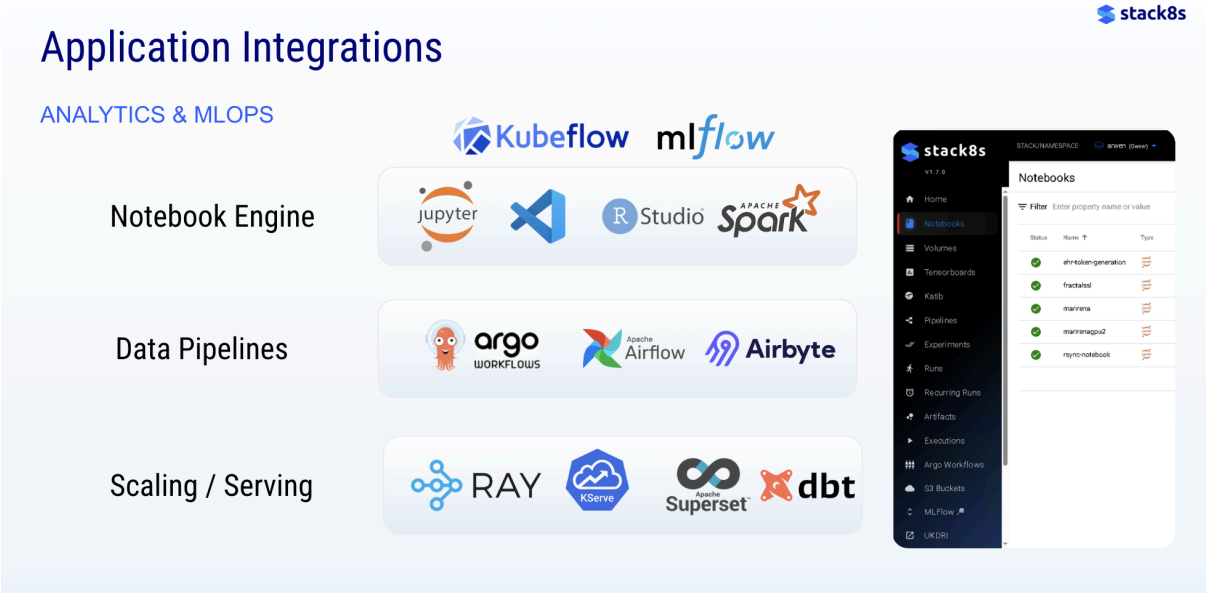
stack8s – Sovereign MLOps Platform

G-Cloud 15 Service Definition

1. Service description

stack8s Sovereign MLOps Platform is a fully managed MLOps and data platform that brings together a curated set of open-source tools (including **Kubeflow**, **Apache Airflow**, **MLflow**, **KServe**, **Ray**, **Feast**, and **Superset**) with integrated Backend-as-a-Service capabilities. It is deployed exclusively on UK and EU sovereign cloud providers to meet strict data residency and regulatory compliance requirements, enabling public sector and regulated organisations to industrialise machine learning without building and operating their own complex infrastructure.



The platform provides end-to-end support for the ML lifecycle: data ingestion and feature engineering, experimentation, training, evaluation, model registry, deployment, monitoring, and ongoing optimisation. Kubeflow underpins notebooks, pipelines, and experiments; Airflow orchestrates production data and ML workflows; MLflow tracks experiments and model artefacts; KServe powers model serving; Ray supports distributed computing; Feast manages feature stores; and Superset provides analytics and dashboards. Backend-as-a-Service components (for example, PostgreSQL, object storage, and APIs) are integrated so teams can build complete ML-powered applications on the same sovereign platform.

stack8s manages all underlying infrastructure, including Kubernetes clusters, storage, networking, security baselines, and multi-cloud deployment across approved UK/EU providers. Customers access the platform through web-based interfaces and RESTful APIs and can focus on ML innovation, data science workflows, and application development rather than infrastructure operations.

CONFIDENTIAL - DO NOT DISTRIBUTE WITHOUT PERMISSION

2. Service features and benefits

2.1 Service features

- UK/EU sovereign hosting with strict data residency controls
  All platform components and customer data are hosted exclusively on UK and EU sovereign cloud providers, with controls to ensure data does not leave agreed jurisdictions. This supports GDPR compliance and departmental policies on data locality for ML and AI workloads.

- Managed Kubeflow ML platform (notebooks, pipelines, experiments)
  The service includes a fully managed Kubeflow environment, providing Jupyter notebooks for experimentation, pipeline authoring and execution, and experiment tracking. Data scientists and ML engineers can design, run, and manage ML workflows through Kubeflow's UI and APIs without administering clusters or storage.

- Apache Airflow for production data workflows and ML orchestration
  Airflow is provided as a managed service for orchestrating ETL, data preparation, and ML workflows in production. Users can define DAGs to coordinate data pipelines and model lifecycle tasks, integrating with Kubeflow and other platform components.

- KServe for model serving (PyTorch, TensorFlow, XGBoost, custom)
  KServe-based model serving allows deployment of models built in frameworks such as PyTorch, TensorFlow, XGBoost, and custom runtimes. It supports CPU and GPU-backed endpoints, autoscaling, and traffic routing for blue/green or canary deployments.

- Open-source BaaS with PostgreSQL, object storage, and APIs
  Backend-as-a-Service capabilities provide managed PostgreSQL, object storage, and APIs so teams can store datasets, features, metadata, and application state within the same sovereign platform. This reduces the need for separate database and storage administration for ML projects.

- Integrated analytics and feature management (Ray, Feast, Superset)
  Ray supports distributed computation for large-scale data processing and training jobs, while Feast provides feature store functionality for consistent feature definitions across training and inference. Superset delivers visual analytics and dashboards for monitoring ML performance, data quality, and platform usage.

- Multi-cloud deployment and resilience
  The platform runs across multiple UK/EU sovereign cloud providers with multi-cloud architecture designed to avoid vendor lock-in and provide geographic redundancy. Automated failover, backup replication, and infrastructure-as-code recovery support high availability and defined RTO/RPO

CONFIDENTIAL - DO NOT DISTRIBUTE WITHOUT PERMISSION

targets.

2.2 Service benefits

- Focus on ML innovation while stack8s manages infrastructure
  Teams can prioritise data science, feature engineering, and model improvement rather than maintaining Kubernetes clusters, workflow schedulers, model servers, and storage. This shortens the path from experimentation to production and reduces the need for in-house platform engineering expertise.

- Meet UK GDPR and compliance requirements with controlled data residency
  By restricting deployment to UK/EU sovereign cloud providers and enforcing strong data residency controls, the platform supports GDPR and public sector data governance expectations. Organisations can demonstrate where data is stored and processed throughout the ML lifecycle.

---

3. Service constraints

- Sovereign-only hosting
  All deployments are hosted exclusively on UK and EU sovereign cloud provider infrastructure; use of other regions or global cloud locations is not supported. This constraint ensures data sovereignty and regulatory compliance but may limit deployment options where non-UK/EU regions would otherwise be used.

- Planned maintenance windows
  Maintenance windows are scheduled during off-peak hours with advance notice to customers to apply security patches, upgrades, and configuration changes. During these periods, short interruptions or degraded performance may occur for some components.

---

4. Technical requirements

- Client access requirements
  Users require a modern web browser (Chrome, Firefox, Safari, or Edge, latest version) to access web-based interfaces for Kubeflow, Airflow, MLflow, Superset, and platform management. A stable internet connection is necessary for interacting with the UI and APIs.

CONFIDENTIAL - DO NOT DISTRIBUTE WITHOUT PERMISSION