

Datavid ROVER Service Definition

Overview

Datavid Rover is a cloud-based semantic data and knowledge platform designed to help public sector organisations turn fragmented structured and unstructured information into trusted, usable knowledge. Rover enables users to discover, connect, and reuse information across departments, systems, and data silos, supporting better operational delivery, research, and decision-making.

Rover builds an AI-ready semantic layer over existing data sources, combining metadata management, taxonomies, ontologies, and knowledge graphs with advanced search and discovery capabilities. It allows organisations to understand not just documents, but the relationships between people, policies, suppliers, datasets, and concepts. This improves findability, transparency, and reuse of information at scale.

The service includes semantic search, graph-based navigation, insights, and private LLM-powered assistance for summarisation and question answering, grounded in the organisation's own data. Rover is API-first, with graph, search, insights, agentic AI, and Model Context Protocol interfaces that allow its capabilities to be embedded into existing digital services and user interfaces.

Rover is delivered as a secure SaaS or managed service, hosted on AWS, with strong governance, auditability, and role-based access controls. It is suitable for discovery, alpha, beta, and live services and helps organisations reduce duplication, accelerate insight, and deliver more effective digital services.

Service Description

Service features

- Data Intelligence Platform
- Taxonomy and Ontology Management
- NLP, machine learning, semantic search, open data
- Semantic Search/Cognitive Search
- Knowledge graph
- Entity Extraction / Entity Classification
- Knowledge Management

- Data Hub or Data Lake
- Agentic AI
- RAG and GraphRAG

Service benefits

- collaborative taxonomy and ontology management
- quickly manage data/content on the move
- Faster time to insight through AI-assisted research
- Transparent knowledge discovery, provenance and auditability
- Relationship intelligence across people, suppliers, and data
- Private, secure LLM-powered knowledge assistant
- Pre-built enrichment pipelines, ontologies, and APIs
- Built-in semantic indexing and cognitive search

Information Assurance

Our team has experience with PROTECT, RESTRICTED, CONFIDENTIAL and SECRET impact level.

Ordering process

Easy to order contact us at info@datavid.com

Invoicing process

Invoice raised after acceptance of each milestone

Termination terms

1 month written notice by consumer.

12 months if service is removed from G-Cloud, allowing enough time for customers to migrate.

Data restoration/service migration

Our DevOps Team will help with migration of service to another provider and provide knowledge transfer to new team.

Consumer responsibilities

Provide description of the use case to be hosted in Cloud.

Product owner of the service or key business stakeholders present for any clarification to the team.

Trial service

Yes, available for 2 weeks with no obligation to the consumer.

Data extraction/removal

The data extraction and removal process for Datavid Rover is designed to be secure, transparent, and aligned with public sector exit and data portability requirements.

At any point during the contract, and at contract end, customers can request data extraction. Data can be exported using Rover APIs to retrieve documents, metadata, entities, relationships, enrichment outputs, and knowledge graph structures in open, machine-readable formats such as JSON, CSV, RDF, or Parquet. For large data volumes, Datavid can provide bulk exports using logical database backups or snapshots.

Where required, Datavid can deliver a complete export to customer-controlled object storage such as AWS S3, Azure Blob Storage, or equivalent BLOB storage. Exports include supporting manifests, schema definitions, and documentation to support onward reuse.

Once the customer confirms successful transfer and validation of exported data, Datavid securely removes remaining customer data from the Rover environment in line with contractual terms, ISO 27001 controls, and regulatory requirements. Data removal is verified and, where required, evidenced to the customer.

Data storage and processing locations

Datavid Rover is hosted on AWS. Customer data is stored and processed in the AWS region agreed at contract call-off. For UK public sector customers this is typically AWS London (eu-west-2) to keep data in the UK, but we can also support EU regions (for example Ireland, eu-west-1) where a buyer's data residency policy requires it. Rover is deployed across multiple AWS availability zones within the chosen region for resilience. Backups, logs, and monitoring data are retained in the same agreed region unless the customer requests a different, compliant location.

Deployment models

1. SaaS (Datavid-managed): Rover delivered as a fully managed service with customer-specific isolation (dedicated Kubernetes cluster per customer).
2. Managed single-tenant in customer cloud: Rover deployed into a buyer-owned AWS/Azure account/VPC, managed by Datavid under an agreed operating model.
3. Hybrid integration: Rover hosted as above while integrating securely with on-premise or private network data sources.
4. Private/isolated environments: Separate non-production and production environments, with promotion controls, for regulated delivery.