



DATA SHARING PLATFORM

Our Data Sharing Platform is a resilient, scalable and cloud-based solution for publishing data. It hosts data for a wide range of our public and private sector clients, and is available as a fully hosted and managed enterprise solution, or elements can be installed locally on your own

infrastructure. It is widely used for publishing linked data on data.gov.uk, including data at: environment.data.gov.uk, location.data.gov.uk, landregistry.data.gov.uk, and many others.

We offer the platform as a fully hosted and managed service for publishing linked data; In addition, we can install customized elements of the platform on a client's own infrastructure.

Our hosting support for the full system includes all regular maintenance, monitoring and backups. We will provide reports to the clients on the usage of the system – the precise details of the data reported will be agreed with the client during the setup phase. We also provide an incident reporting service. Additional support options are available at extra cost.

The production system we run for the Environment Agency hosts four live services (Bathing Water quality, Flood warnings and river levels, Water Quality Archive, electronic Public Registers) which have averaged 99.998% availability for the last year. As the system is fully replicated routine maintenance can be carried out without taking the system off-line; there is no need for scheduled maintenance windows when the system is out of service.

BENEFITS



Robust, reliable service for publication of sustainable, trusted and usable 5-star data



Deployment flexibility on the cloud or within your own infrastructure



Adaptive and flexible platform that can be grown to meet your changing needs



Provide your partners or customers with integrated, live updates twenty-four hours a day

FEATURES

- Fully standards-compliant linked data publication platform
- Widely used within the UK public sector
- Replicated for fault-tolerance and scalability
- Entry level service also available

ARCHITECTURE

The platform architecture has 3 main tiers: load balancing and routing, application services and storage.

Optionally we can provide additional upload mechanisms which will integrate with client's existing workflows to support "business as usual" publication of linked data.

The platform is customisable and can also host applications running on top of the data.

TECHNICAL DESCRIPTION

The Epimorphics Linked Data Platform is used by the Environment Agency, Land Registry and commercial customers for linked data publication.

It is comprised of:

- - A Linked Data API engine, provided by Epimorphics' implementation of the LDA (ELDA) or specialised alternatives when needed
- - Text search, provided by Apache Lucene
- - A fully compliant SPARQL 1.1 endpoint, provided by Apache Jena ARQ
- - A scale-out, fault-tolerant runtime platform, hosted in Amazon Web Services
- - An update controller for managing coordinated updates to replicated services
- - Support for converting data from source formats such as CSV to linked data for publication
- - Support monitoring and scaling the underlying cloud infrastructure