

## Data in-motion Visibility & Accessibility

Lenses, a DataOps platform, accelerates time to value, opening up data streams to a wide audience. Lenses enables rapid construction and deployment of data pipelines at scale, for enterprises, with governance and security.

DataOps is currently transforming data management. Building a data-driven culture mandates that all data personas work with data, enabling participation from the entire business. Organisations are trying hard to expose their data via platform teams. As a result, many end up with customised DIY solutions with in-house engineering teams spending most of their time building infrastructure and tooling. Data should be in the hands of its users, as simple as their email, to enable innovation and minimise time to value. Lenses empowers this.

Get more from you data with Lenses and DataOps, leverage your data and existing data skills to iterate quickly and move to production faster.

### Focus on:

- ➡ Leveraging your whole organisation
- ➡ Data pipelines over infrastructure
- ➡ Configuration over code
- ➡ Data analytics over engineering
- ➡ Operations made easy

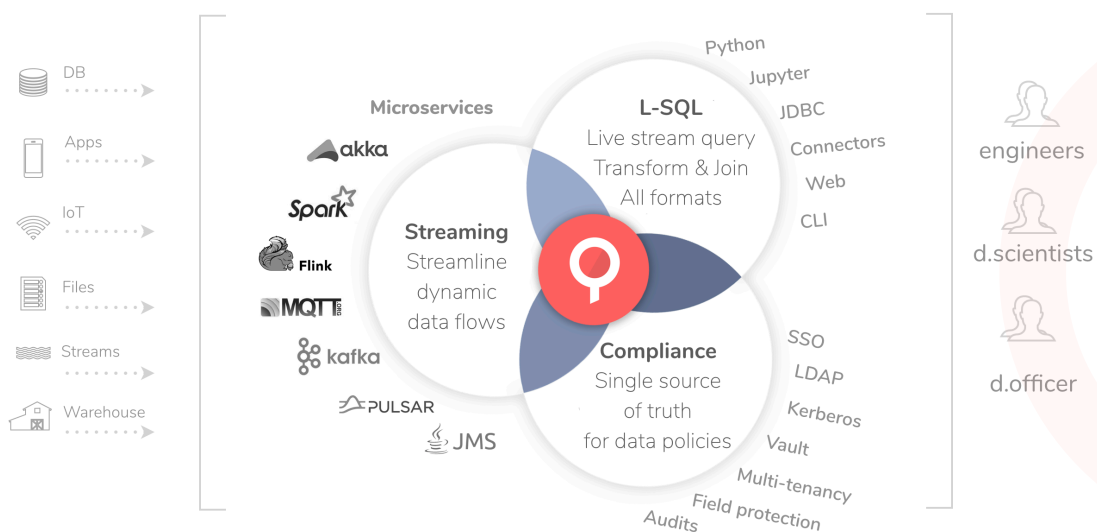
Development to Data with Operations.

### BENEFITS

- ✓ Faster time to value
- ✓ Empowerment of data literate business users
- ✓ Reduced engineering effort
- ✓ Focus on data and business requirements

### APACHE KAFKA® ENABLER

- ✓ Access live and batch data
- ✓ Power your teams with access to real time streams
- ✓ Move to production faster with a comprehensive to access data



# SQL, the language of data

## Why limit visibility and innovation only to specialists?

Lenses provides state of the art SQL visibility into your data flows, engaging all areas of your business as you accelerate towards data driven platforms. Lenses SQL engine sits above your middleware, including Apache Kafka®, allowing you to focus on discovering and innovating with the data you know best.

## Browsing and Continuous queries?

Lenses SQL supports both browsing and continuous queries to join, filter and aggregate data. Integrate the Lenses SQL engine into your applications with our Python, Redux or Golang client. You can even hook up your Spark or Flink jobs with the JDBC driver.

Feature highlights:

- ✓ Data browsing with built-in protection via multiple channels
- ✓ Continuous queries for enrichments, transformations, filtering, joins
- ✓ All major math, string, date and aggregates functions
- ✓ Access and manage metadata & schemas: header, key, value, timestamp and partition of the messages

## Any format, any data?

We know your data isn't always nicely presented to you in the format you want, that's why Lenses supports all the major formats, including Protobuf for those of you concerned about performance!

Support for:

- ✓ CSV, XML, JSON, Avro, Protobuf, Primitives Types, Pluggable custom formats

The screenshot displays the Lenses SQL interface. On the left, a sidebar lists topics and their schemas. The main panel shows a SQL query: `SELECT MMSI, Speed FROM position_reports WHERE S LIMIT 1000`. Below the query, the 'Run Query' button is visible. The results section shows a table with columns: Second, Speed, Status, sign, sin, sinh, sqrt, strcasecmp, sub, substr, substring, sum, swapcase. The results are fetched in batches of 1,000. A detailed view of a result is shown, including fields like location, Radio, Latitude, Longitude, Speed, Turn, Repeat, MMSI, Course, Timestamp, Heading, Maneuver, Status, and Type. On the right, a 'Processors' panel shows a data flow diagram with nodes for 'TOPIC', 'TABLE', 'FILTER', 'JOIN', 'GROUPBY', 'AGGREGATE', and 'TOPIC'. Below the diagram, a 'Topic Details' panel shows the schema for 'cc\_payments\_stream'.

## Control Access

**Authentication & Authorisation?** Lenses makes data access simple and is a gateway for innovation for your teams. At the same we make sure that you have enterprise ready integrations. Lenses provides Role Based Access Control for basic authentication, but also LDAP and Kerberos integrations. Multi-tenancy is simplified, with whitelisting and blacklisting that also applies to any applications trying to access data. If you are white/blacklisted you can't see any application using the data or create any that can access it.

Support for:

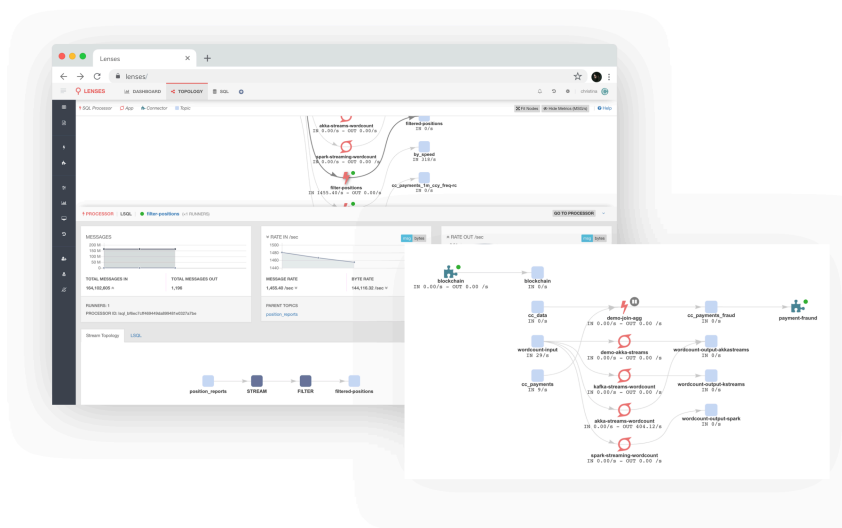
✓ ACLs, Quotas, RBAC, White and blacklisting, TLS support

**Field Policies & Compliance & GDPR?** Do you know where your data is going? Which systems in your data platform contain PII data, which topics and which applications are using them? Lenses applies policies on the data flowing through your platform enabling you to answer these questions and track sensitive data through your organisation.

**Governance?** Building on the enterprise security features Lenses goes a step further, all actions are audited, including queries, giving you governance and lineage over "who did what and when" on your platform.

**Provenance & Lineage?** Lenses tracks and monitors underlying middleware, such as Apache Kafka® and your application landscape, *the business value*. Lenses provides a dynamic, real time topology view of your application landscape providing lineage. Promote flows between environments by plugging the Lenses CLI into your continuous delivery pipelines and store your topologies in version control. Compare your desired state against the realised and eliminate drift.

- ✓ Real time view of Application topologies
- ✓ Dynamically built from actual realised state
- ✓ Kubernetes support
- ✓ Drill down to applications, topics and internal topologies
- ✓ Custom application support for KStreams, Akka Streams, Spark, Flink and JVM apps
- ✓ Identify data stores holding GDPR data.



## Lenses Native Clients



Lenses provides a wide ecosystem of clients to access and enhance data information. Python, Javascript and Golang clients are available and additionally a JDBC driver. Integrating via Lenses clients helps you run queries and bring data to your own tools backed by the security controls and policies. This way, you get out of the box lineage, data redaction features and a central secure gateway to access your data.



## Features on Data Visibility & Accessibility

Manage Data	
Lenses SQL Engine	Transform data using SQL, regular expressions, nested data operators, etc <a href="#">View more</a>
Supported Formats	Primitive Types, JSON, XML, CSV, Avro, Protobuf, Custom Serializers
Complex Types	Lenses SQL support for complex types, arrays, unions, structs, maps
Data Preparation	Apply transformations, calculations, and aggregations
Insert Data	Insert data into topics with payload suggestions .i.e. Testing event driven microservices
Delete Data	Delete data from compacted topics
Browse Data Clients	
Web UI	Subscribe & Query Live or Retained Data Streams with SQL or quick filter Navigate to Partitions, offsets, Access metadata, Tree, Grid, Raw Data viewer, Download selected data set
Endpoints	Rest / Websocket
CLI	Lenses Command Line Interface to automate running queries
JDBC Driver	JDBC driver plugin to access data in motion for your custom apps or BI tools <a href="#">View here</a>
Native Clients	Python client for Jupyter notebooks and Golang client for mirco-services. Redux middleware for building real time dashboards on live queries
Continuous Queries	Scalable continuous processors in Kubernetes
Query Management	View and Manage live queries from users via the admin CLI
Access Metadata	Message metadata (key value types, headers, schemas) Schemas (integrate with schema registries)
Control Access	
Role Based Access	Lenses supports role based authorisation to protect your cluster and your data, especially in a multi-tenant environment. Also provides a no-data role for operation and support teams
Basic Authentication	Token based authentication scenarios
LDAP	Integrate via memberOf or pluggable LDAP, or AD (Active Directory) and provide the people in your organization role based access to Kafka streaming data
Kerberos	JaaS, Utilise Kerberos Authentication, Authorization, Single sign-on
TLS Certificates	Authenticate and Authorise Clients and Application via signed TLS certificates
ACLs	Manage Access Control Lists (Kafka ACLs) via Lenses interface
Multi-tenancy	Blacklist and Whitelist capabilities
Governance & Lineage	
Audit Logs	Immutable audit logs for all user actions & data queries
Data Schemas	Schema Registry + Lenses Metastore
Topologies	Application / Topology level view of data-pipelines, highlights relations of data with private or GDPR regulated data

## Find out more

- ➡ [How to Access data in-motion with Lenses](#)
- ➡ [Lenses SQL Engine for Browsing & Processing Data](#)
- ➡ [Lenses Resources](#)

## About Us

A team of passionate distributed system specialists based in the City of London with high-frequency trading, investment banking, betting, media and retail experience are the creator of **Lenses** ®

We aim to define Data Ops and assist companies open up their business data to all relevant users seamlessly by giving them the power of data operations. DataOps is currently transforming data management. Building a data-driven culture mandates that all data personas work with data, enabling participation from the entire business.

Modern organisations are trying hard to expose their data via platform teams. As a result, many end up with customised DIY solutions with in-house engineering teams spending most of their time building infrastructure and tooling.

We believe that data should be in the hands of its users, as simple as their email, to enable innovation and minimise time to value and used ethically.

Wanna give it a try ?

**[Lenses.io](#)**