DATA SHEET

DEC 2018 Lenses.io

Data Transformation

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



LensesSQL Processors

Lenses, with its large collection connectors, enables simple and intuitive collection and storage of all types of data streams, from Bloomberg data to MQTT. However, to extract value and drive business insights data often has to be transformed, filtered, join and aggregated. Lenses SQL engine, in addition to SQL Browsing and Continuous queries, provides Continuous Processors.

Continuous Processors, using a unified SQL syntax, allows for rapid creation, deployment and monitoring of applications with the ability to store the results back to a middleware. Lenses provides three deployment models out of the box, In-Process, Apache Kafka® Connect and Kubernetes. Lenses manages the deployment and monitoring allowing you to focus on becoming data-driven.

In-Process, a development mode, allows teams to quickly literate and try out their ideas. The Lenses development box, a docker with the complete stack, enables end to end flows to be created locally on a laptop. For enterprise workloads, we recommend to use the Connect or Kubernetes mode.

Connect mode provides a fault tolerance and scalable platform for the enterprise user to deploy and manage Continuous Processors. Simply add the specialist Lenses SQL Processor connector to your Apache Kafka® Connect clusters and Lenses will do the rest. Scale and achieve fault tolerance with a vanilla Apache Kafka® distribution. No extra services required.

Kubernetes mode provides an ideal platform to scale processing and deploy microservices. Lenses can deploy, monitor and manage Continuous Processors in Kubernetes. Lenses providers a Docker image and Helm charts. Processors can be deployed outside of Lenses and will still track, monitor and manage them.

CALENSES	SQL DOLICIES			
Topics System Topics pos militared-positions militared-positions2	SELECT MASI. Speed FROM position_reports INNER LIMET 1000 Run Query Results (1000) Recent Queries (1) O STATS Fetched: 1,000	E S Second Second Status sign sin sin sin sin sin sin sith sgrt stap sin sith sgrt stap sin sith sgrt stap sin stap	column column column fra fra fra fra fra	
titlered-positions3 position_reports	► KEY: { MMSI: 265627610 } ► VALUE: { Second: 28, location: 56.040950,14.583028, l OFFSet: 1628254515 PARTITION: 1 TIMESTAMP: 2018-09-29 l	sun swapcase Rad	fn fn 	
_key.MMSI key Timestamp Maneuver Turn Status Redio RAIM Second Haadng Course Longitude Accuracy Speed MMSI Repat Type	■ KKY: (MASI: 265391570) ● VALUE: General: 28, Genera: 28, Genera: 28, Genera: 28, Ge	Radio: 66591, Latitude: 59.32128333 07.41.46 words]: 1 presempte flad detacle: 0 words]: 1 presempte flad detacle: 0 words in words to get a set of the	S33334, Longitude: 18.10344, Spee	OCUTY ACCELST TO CONTRACT, AND

Custom Applications

If you microservice can't be written in SQL Lenses can still help you monitor and govern you applications. Lenses comes with client topology and metrics libraries that can be included in your applications. This allows Lenses to monitor and track your custom applications including

- ✓ Kafka Streams
- ✓ Akka Streams
- ✓ Apache Spark
- ✓ Apache Flink



Monitor & Govern

Lenses sees it all, from infrastructure status to the application landscape. Lenses builds a live topology of your application landscape, complete with real time status and metric updates for connectors, topics, continuous processors and custom applications. Each action, including queries, is audited.

Features on Data Transformation

Data Analysis & Preparation		
Know your data	Discover, Explore and Query your Data o understand how events and signals from are captured	
Access Data	Provided you have authorization to access the data, easily access them irrespective of the underlying data format or serialization protocol	
Structure & Organise Data	Organize and structure your data in efficient data formats	
Cleanse Data	Filter, map and prepare the data for your business domain	
Transform Data	Apply the necessary transformations and aggregations and enrich the data on the fly	
Move Data	Move data between applications and source and target data systems	
Extract Insights	Perform fast SQL operations for instant response to BI dashboards, reports, and machine learning applications. Lenses SQL integrates with native clients like JDBC driver, ReduxJS middleware for custom frontend apps, Python for integration with Jupyter or other notebooks.	
Analytics	Continuous scalable queries for real-time filtering, enrichment and data transformation. Perform counts and anomaly detection use cases with a simple query. Output data to multiple systems and generate live dashboards to meet your SLAs	
Distributed Joins	Express massively parallel distributed joins, without having to worry about partitioning semantics	
Self-service Infrastructure	Move your SQL logic into production in seconds, and enjoy monitoring alerting and control over your data pipelines	
Anonymization	Anonymize and Obfuscate data while it transitions across your staging area	
Field Protection	Apply field level protection for sensitive and PII data	
Data Transformation & Stream Processing		
SQL Processors	SQL processors are continuously unbounded SQL queries, specifically for streaming data. Each query is providing its internal topology which is visualised in Lenses UI as interactive nodes, where we have live input and output data.	
SQL Processors Scalability	 SQL processors provide 3 execution modes: In process: which is good for development purposes Kubernetes: where each runner is managed by a Kubernetes pod Connect: where each runner is orchestrated by Kafka connect workers 	
Native scale via Connect Workers	Distribute your stream-processing jobs to Kafka-Connect clusters. Lenses SQL stream	
	processors provide a native integration with Connect workers to scale up or down. The number of processor runners can simply be specified via Lenses UI.	
Native scale via Kubernetes	processors provide a native integration with Connect workers to scale up or down. The number of processor runners can simply be specified via Lenses UI. Run and manage your data stream processing application within Kubernetes. Lenses SQL stream processors provide a native integration with Kubernetes to scale up or down. The number of processor runners can simply be specified via Lenses UI. Lenses manages the pod lifecycle and exposes the logs for debugging purposes.	
Native scale via Kubernetes SQL Processors Capabilities	 processors provide a native integration with Connect workers to scale up or down. The number of processor runners can simply be specified via Lenses UI. Run and manage your data stream processing application within Kubernetes. Lenses SQL stream processors provide a native integration with Kubernetes to scale up or down. The number of processor runners can simply be specified via Lenses UI. Lenses manages the pod lifecycle and exposes the logs for debugging purposes. Self-service data preparation Topic to Topic transformations Filter and Enrich data Join the stream (with repartition support) Aggregations 	

Monitoring	For each processing entity (SQL or Custom) Lenses collects performance metrics for which threshold alerts can be applied. For example data in/out, rates, data size.
Alerting	Threshold alerts can be specified for the input data for each SQL processor

Data Pipelines / Streaming ETL ready		
Global Topology	Lenses identifies the related entities (topics, connectors, processors) and constructs the pipeline for each flow. Topology has interactive nodes and	
Lenses Source Connectors	Bring data onto Kafka from multiple sources: Twitter, BlockChain, MQTT, JMS, RDBMS, FTP, File location, Yahoo, Bloomberg, Blockchain, Cassandra, CoaP, ReThink	
Lenses Sink Connectors	Popular distributed NoSQL and caching systems: Cassandra, DocumentDB, Elastic Search, SolR, HDFS, RDBMS, HazelCast, InfluxDB, MongoDB, HBase, Kudu, Redis, Pulsar, VoltDB, Rethink and others	
SQL for Ingestion	Lenses connectors (see github.com/Landoop/StreamReactor) support SQL configuration to give extra capabilities and integrate with target systems	
Custom Connectors	Lenses is compatible with all Kafka connectors as it implements the connect API. Lenses makes connectors available for spin up instances as well as available in the global topology.	
Change Data Capture	Extract Change Data Capture (CDC) events from MainFrame or RDBMS via bin-logs.	
SQL Stream Processing	SQL processors can be part of your ETL Flow (see Data Transformation)	
Custom Stream Processing	Hook custom applications to your flows. Native support for Spark, Kafka Streams and Akka Streams and compatible with other types. (see Data Transformation)	
Manage Lifecycle	Admin actions for Connectors & Processors, to edit, scale, start/stop.	
Import / Exports	Lenses CLI supports import/export of different elements as configuration, to create repeatable flows and promote across environments.	
Multi-tenancy	Lenses provides multitenancy support via blacklisting/whitelisting which can protect the relevant flows via topic-centric configuration. Restricts the access to involved connectors, processors and consumers.	
Alerting	Apply data consumption threshold alerts for each entity involved: consumers, sink connectors, SQL processors	

More on Data Access

- 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 datadriven 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

All rights reserved. Lenses is trademark and registered trademarks of Landoop Ltd. in the United Kingdom and other countries. All other brand names, product names, or trademarks belong to their respective owners.