

Oracle Stream Analytics



Oracle Stream Analytics allows users to process and analyze large scale real-time information by using sophisticated correlation patterns, enrichment, and machine learning. It offers real-time actionable business insight on streaming data and automates action to drive today's agile businesses.

From its interactive designer users can explore real-time data through live charts, maps, visualizations, and graphically build streaming pipelines without any hand coding. These pipelines execute in a scalable and highly available clustered Big Data environment utilizing Spark integrated with Oracle's Continuous Query Engine to address critical real-time use cases of modern enterprises.

The screenshot displays the Oracle Stream Analytics interface for a pipeline named 'RT Marketing'. The pipeline consists of several stages: CustomerActivity, Customers, Enrich, Near Mail, In Mail, Predict Sales, Likely Buyers, and Send Offers. The 'Enrich' stage is currently selected, showing its configuration. The 'Enrich' stage has two sources: 'CustomerActivity' and 'Customers'. The 'Correlation Conditions' section shows 'CUSTID equals custid'. Below the pipeline, the 'Enrich - Live Output' table is visible, showing real-time data rows.

custid	prodid	lat	long	timestamp_1	FCUSTID	AVGDISCOUNT	CUSTID	NAME	STORESALES	WEBSALES
65	5000	37.326581318731236	-121.94350414255631	2018/02/23 16:06:21	65.0	5.0	65	Paul Stone	79.0	492.0
64	5000	37.31907370411125	-121.95892823751905	2018/02/23 16:06:21	64.0	38.0	64	Peter Pyle	443.0	828.0
63	5000	37.33933722701945	-121.9424720157426	2018/02/23 16:06:21	63.0	16.0	63	Paul Frey	440.0	861.0
62	5000	37.3219963022251	-121.9518819191324	2018/02/23 16:06:21	62.0	7.0	62	Thomas Waters	839.0	651.0
61	5000	37.3277369502072	-121.94353515670371	2018/02/23 16:06:20	61.0	2.0	61	Robert Storm	127.0	921.0
60	5000	37.322164637156478	-121.93491954302186	2018/02/23 16:06:20	60.0	22.0	6E+1	Linda Martel	947.0	1034.0
59	5000	37.32039903200034	-121.94915530229778	2018/02/23 16:06:20	59.0	33.0	59	Chris Martinez	795.0	743.0

Designing Streaming Pipelines with Oracle Stream Analytics

About Oracle Stream Analytics

The quantity and speed of both raw infrastructure and business events is exponentially growing in IT environments. Whether it is real-time marketing and dynamic pricing for retailers, streaming satellite data for the military or real-time vehicle location data for transportation and logistics businesses, companies in multiple industries must handle large volumes of complex data in real-time. In addition, the explosion of mobile devices and the ubiquity of high-speed connectivity add to the explosion of mobile data. At the same time, demand for business process agility and execution has also grown. High-velocity data brings high value, especially to volatile business processes. However,

KEY FEATURES AND BENEFITS

- Interactive Design of Stream Analytic Pipelines through Graphical Editor
- Reusable Business Solution Patterns available for users to select a familiar solution analysis
- Graphical visualization of real-time data in various chart styles, maps and tabular displays
- Analyze and correlate geospatial information in streams, graphically define and introspect location data and rules
- Predictive Analysis based on a wide selection of Machine-Learning models, integration with industry-standard PMML format
- Change data capture through integration with Oracle GoldenGate, analysis and transformation of captured operational mission critical transactions.
- Rich Set of Event Stream sources and targets such as the streaming of CSV, JSON, and Avro files in Kafka and JMS queues
- Robustness, Scalability and Speed through parallel, in-memory processing in Spark Big Data environment.


some of this data loses its operational value in a short time frame. Actionable insights need to be detected immediately as new data is available.

These two trends have put pressure on organizations to increase their capability to support event-driven architecture patterns of implementation. Real-time event processing requires both the infrastructure and the agile design environment to execute on event processing requirements.

Common Use Cases

The Oracle Stream Analytics platform targets a wealth of industries and functional areas. The following are typical use cases:

- **Marketing and Retail:** Real-time marketing to create advertisement and offers tailored to geographical and situational context of customer, dynamic pricing based on situational factors.
- **Supply Chain and Logistics:** Ability to track shipments in real-time and detect and report on potential delays in arrival. Control stocking levels based on change in demand and shipping predictions.
- **IT Systems:** Monitor and recognize failures and performance bottlenecks in applications or servers and trigger corrective measures. Identify security risks and incidents and initiate response.
- **Telecommunications:** Ability to perform real-time call detail (CDR) record monitoring and distributed denial of service attack detection.
- **Financial Services:** Perform real-time risk analysis, monitoring and reporting of financial securities trading and calculate foreign exchange prices.
- **Transportation:** Ability to create passenger alerts and detect baggage location in case of flight discrepancies due to local or destination-city weather, ground crew operations, airport security, etc.
- **Public Sector/Military:** Detect dispersed geographical enemy information, abstract it, and decipher high probability of enemy attack. Ability to alert the most appropriate resources to respond to an emergency.
- **Insurance:** Ability to learn and to detect potentially fraudulent claims.



Define Geospatial processing in Oracle Stream Analytics

Oracle Stream Analytics allows users to process and analyze a large scale of real-time information by using sophisticated correlation patterns, enrichment, and machine learning, and provide insights and real-time decisions to drive today's agile businesses.

RELATED PRODUCTS

The following products complement Oracle Stream Analytics to process information in real-time:

- Oracle Big Data Cloud
- Oracle Event Hub Cloud Service
- Oracle GoldenGate
- Oracle Integration Cloud
- Oracle Data Integration Platform Cloud
- Oracle Data Integrator
- Oracle IoT Cloud
- Oracle Analytics Cloud


RELATED SERVICES

The following services support Oracle Main Product:

- Update Subscription Services
- Product Support Services
- Professional Services

Rich Set of Streaming Patterns


Oracle stream analytics provides an extensive pattern library to analyze complex event patterns for a wide array of use cases such as Real-Time Log Analytics, Digital Marketing, Predictive Maintenance, Fraud Detection, Smart Inventory, Dynamic Pricing, etc. For example patterns can detect correlated events, such as a door closing within 10 seconds of opening it, spikes in values, such as a temperature spike of 30% outside of the average within a given time, or detecting a W-pattern in financial trading. Patterns can also detect if a weight of manufactured items is slowly trending lower or the reading falls outside acceptable norms. This signals a potential problem or future maintenance need.



Oracle Stream Analytics Pattern Library

Interactive Designer UI

Oracle Stream Analytics is democratizing the ability to analyze streams. It provides an interactive UI to explore streaming data through visualizations or live data views and build streaming pipelines without coding. Users can introspect data from streams and get immediate feedback on how patterns applied on the live data create actionable results. Streaming pipelines can be built through a graphical designer that shows a diagram of stages that incoming data runs through. Data can reviewed and manipulated in a spreadsheet-like tabular view, allowing users to add, remove, rename, or filter columns to obtain the desired result.



Tabular and Chart Visualizations of Streaming Data

Location and Geo-Spatial Capabilities

Geo-Spatial analysis is often a component of event processing as location is a typical form of ephemeral information that needs to be acted upon in real-time. Oracle Stream Analytics provides functionality to build location-aware applications and detect events such as entry and exit from observed areas as well as proximity and speed of traced objects.

Map views allow users to graphically define areas of interest through defining geofences and track locations of objects over time. Maps can be color-coded based on statistical data and visualize proportions based on size.

Predictive Analysis and Machine Learning

While traditional Analytics and Business Intelligence provide results for human decision-making, Streaming Analytics has to make autonomous real-time decisions based on the data present in order to send alerts or to take remedial actions. Scoring and predictions based on machine learning models allows to decide on an action, for example by predicting probability of failure for an observed device and initiating service activity. Oracle Stream Analytics allows to import machine learning models such as classifiers, regression models, clustering models, neural networks, decision trees, and other models using the PMML industry standard. This allows reuse of models created by data scientists using common ML tools within streaming pipelines without custom work necessary.

Change Data Capture with Oracle GoldenGate Integration

Oracle Stream Analytics is integrated with the industry-leading Oracle GoldenGate change data capture platform to process live transaction feeds from transactional sources such as OLTP databases to detect pattern in real-time and prepare and enrich data for analytical stores. Inserts, changes and deletions in the data can be processed, replicated in-memory, or forwarded to the target store. The scalable run-time environment based on Spark allows to match the high-performance throughput of GoldenGate and provide a large scale end-to-end processing solution of streaming data.

Robustness, Speed, and Scalability

Event streams often scale to millions of events per second with stringent requirements of exactly-once accuracy. Oracle Stream Analytics is built on the industry-leading Spark-Streaming standard to allow horizontal scaling in a cluster using on-premise or cloud big data architecture. Spark provides high-speed parallel in-memory processing to scale to any size load, and provide high-availability through redundancy across the cluster. Oracle Stream Analytics combines the standard Spark framework with its CQL event processing framework to add patented event processing with exactly-once delivery guarantees. Distributed caching is provided through the Oracle Coherence In-memory cache that is seamlessly embedded as part of the Spark deployment.



CONTACT US

For more information about Oracle Stream Analytics, visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US



Integrated Cloud Applications & Platform Services

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0318



Oracle is committed to developing practices and products that help protect the environment