

ORACLE CDC SOURCE CONNECTOR

Unlock Valuable CDC Data from your Oracle Database in Real Time with Confluent

Confluent's Oracle CDC Source Connector enables you to:



Avoid up to \$2.5M in prohibitive licensing costs from capturing high-value Oracle DB change events with legacy vendors (e.g., Oracle GoldenGate)



Modernize your databases and data warehouses by streaming valuable operational data to modern, cloud-native technologies to enable real-time use cases & applications



Save ~12-24 engineering months & eliminate operational burden by leveraging an expert-built, fully managed connector for highly complex CDC integrations

Confluent's Oracle CDC Source Connector offers a cost-effective, easy, and reliable way to stream real-time changes from your high value Oracle databases to build modern, cloud-based experiences

To succeed in today's digital-first world, businesses must deliver rich customer experiences and data-driven, backend operations. This requires the ability to react, respond, and adapt to a continuous, ever-changing flow of data from across your organization in real time. However, for many businesses, much of that data still sits in legacy databases like Oracle, which were designed for static, on-premises deployments and not built with the agility expected of today's cloud-based platforms. There is an essential need to stream this data to multiple destinations like analytics platforms, cloud-native databases, and data warehouses, so organizations can innovate and build applications and operations that leverage the data in real time.

As one of the most popular legacy databases, Oracle databases often hold highly critical enterprise data like transactions, customer data, contracts, and other valuable sources of business intelligence. Oracle databases were designed for storing data at rest, hence they struggle to implement continuous real-time syncs and underlying data changes to other downstream systems, such as data lakes & warehouses, BI tools, and other apps relying on real-time alerts and notifications. Change Data Capture (CDC) seeks to solve this challenge by efficiently identifying and capturing data that has been added to, updated in, or removed from Oracle relational tables.

Apache Kafka[®], used by over 70% of the Fortune 500, has been selected as the de facto real-time bridge, acting as a central nervous system to harness the flow of data across all key pieces of an organization in real time: between applications, databases, SaaS layers, and both legacy and cloud ecosystems to build modern digital experiences and operations. To truly tap into Kafka though, many companies look to Confluent's cloud-native and complete data streaming platform, available everywhere their apps & data reside, enabling organizations to accelerate time-to-value while reducing their Total Cost of Ownership (TCO) by up to 60%. Confluent helps businesses unlock this change data to enhance their real-time use cases, which requires bridging legacy technologies to their modern, cloud-based data systems and applications.

The Challenge

Sending Oracle CDC data to Kafka often adds significant complexity and cost for organizations, as few tools exist in the market today to address this specific need at scale. Businesses usually take one of two approaches, each with its own challenges. They can either leverage a third-party vendor where available or look to develop their own Kafka connector for Oracle CDC in-house.

Existing integration solutions for capturing events via Oracle CDC include Oracle GoldenGate for Big Data, which requires purchasing expensive licenses that cost upwards of millions of dollars, not to mention the ongoing support & maintenance costs each year. An additional challenge with this approach is that now you have to manage two different systems for capturing events across your data infrastructure stack - both Kafka and Oracle GoldenGate.

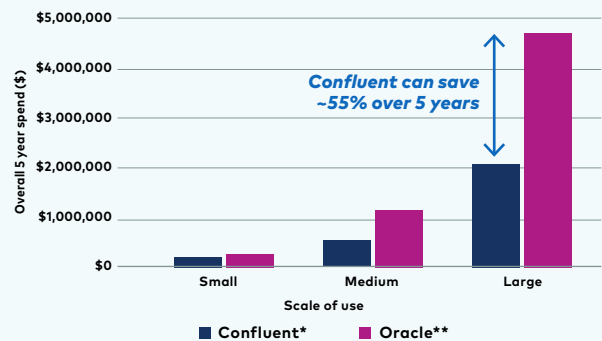
The other option is to develop an Oracle CDC connector internally, which may avoid prohibitive licensing costs, but instead requires significant engineering resources to develop and maintain complex CDC integrations. Even for organizations with significant Kafka expertise, it often takes ~12-24 engineering months to develop an Oracle CDC connector due to ensuring compatibility with different Oracle database versions and configurations, along with continuous mining capabilities. Not to mention the extra cycles spent supporting and maintaining the connector over time, creating additional sources of costly technical debt.

The Solution

Confluent's Oracle CDC Source Connector allows you to cost-effectively, reliably, and securely implement continuous real-time syncs by streaming high-value data from Oracle databases to Confluent, enabling developers to capture change events from each table into a separate Kafka topic.

Compared to Oracle GoldenGate, using Confluent can reduce licensing costs by ~\$2.5M over a five year period for large scale use cases. Using Confluent's pre-built, fully managed connector also frees precious development and operations resources to focus on business differentiating efforts, rather than spending valuable cycles building and maintaining low level data infrastructure tooling.

Confluent vs. Oracle GoldenGate 5 year spend comparison



* Confluent spend is projected Confluent Cloud consumption costs
 ** Oracle spend is Oracle GoldenGate for Big Data perpetual licenses + support & maintenance costs

Confluent's Oracle CDC Connector at-a-glance:

- Supports Oracle DB versions: 11g, 12c, 18c, 19c
- Table Partition Snapshots
- Supports multiple data types – BLOB/CLOB/NCLOB
- Supports Real Application Clusters (RAC)
- Removes need for Oracle GoldenGate for Big Data license
- Single Message Transform (SMTs)
- Connect Logs Events
- Connect Data Preview

Additionally, Confluent's Oracle CDC Source Connector offers superior snapshotting & CDC performance, supports BLOB/CLOB/NCLOB data types, and provides seamless integration with Schema Registry. Confluent also offers over 120 connectors to integrate with non-Oracle environments. Combining this with Confluent's stream processing capabilities in ksqlDB and Kafka Streams enables customers to join, enrich, and transform these data streams in-flight to ensure data is consistent and usable in real time across all of their downstream systems and applications.

Confluent's Oracle CDC Source Connector is available both fully managed in Confluent Cloud or self-managed with Confluent Platform, providing you the ultimate freedom of choice across everywhere your Oracle databases, data systems, and applications reside.

IT use cases:

- Database modernization
- Data warehouse modernization
- Data lake ingestion
- Streaming ETL

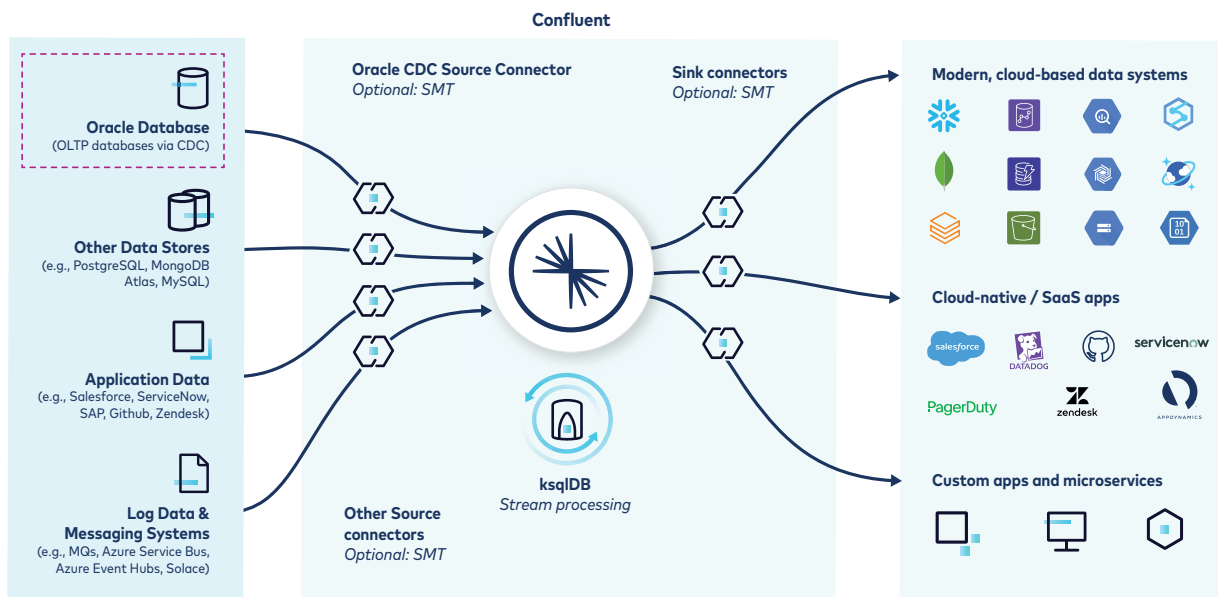
- Data synchronization
- Application integration
- Real-time analytics

Business use cases:

- Customer 360
- Insurance claims processing
- Real-time inventory
- Real-time logistics management

Example Architecture

Unlocking valuable Oracle CDC data to integrate with modern data systems and applications



Our Customers



Getting started is easy

Confluent's Oracle CDC Source Connector is available everywhere your data and apps reside:

Confluent Cloud: <https://docs.confluent.io/cloud/current/connectors/cc-oracle-cdc-source/>

Confluent Platform: <https://www.confluent.io/hub/confluentinc/kafka-connect-oracle-cdc>

Sign up for Confluent and start streaming in minutes.

Visit confluent.io/get-started for more details.

Visit confluent.io/contact to contact us.