

ORACLE

# GoldenGate for Distributed Applications and Analytics 23ai

---

Powering real-time data for modern applications and analytics

June 2024, Version 2.0

Copyright © 2024, Oracle and/or its affiliates

Public

## Introduction

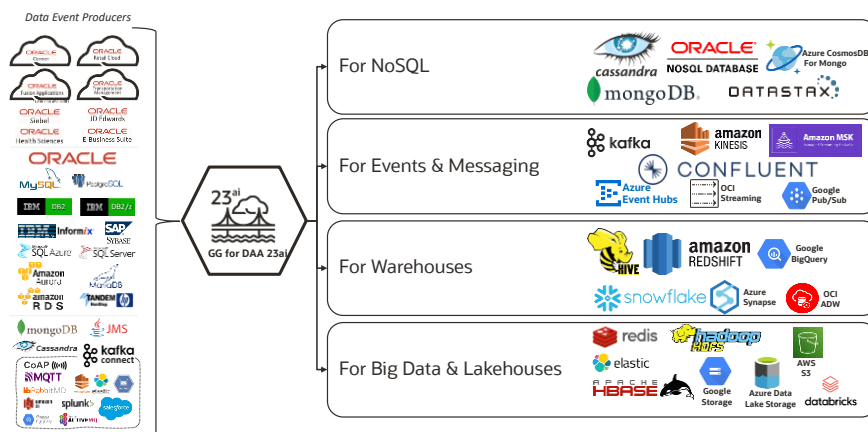
In today's event-driven landscape, microservices architectures are a cornerstone for delivering exceptional customer experiences. This complex landscape requires a reliable real-time data management layer that fuels both analytics services and microservice applications. Traditional batch/bulk ETL data pipelines struggle to keep pace with the high-throughput, low-latency and transaction consistency demands of microservices.

GoldenGate for Distributed Applications and Analytics 23ai (GG for DAA 23ai) is a solution tailored for simplifying the challenges of microservices. This single, user-friendly package empowers data engineers and application developers by reducing integration complexity and streamlining data flow specifically for microservice architectures. GG for DAA 23ai brings 3 major components together:

- **GG for DAA Handlers** are native source and target connectors for message streaming, data/ delta lake, cloud warehouse and NoSQL database technologies. GG for DAA handlers provide low-impact capture and real-time data ingestion capability with high data accuracy and throughput.
- **GoldenGate Data Streams** is a REST interface to your DML and DDL data events exposed as JSON formatted CloudEvents. Data Streams are accessible through the [AsyncAPI](#) standard and channels to enable popular programming languages to directly consume data events.
- **Transaction Manager for Microservices** (MicroTx) protects against multi-datastore inconsistency when microservice application require distributed transaction management.

## GG for DAA Handlers

GG for DAA handlers are native connectors which enable users to ingest real-time messages and capture data from GoldenGate trail files and directly from messaging diverse set of technologies. GG for DAA 23ai handlers support various business cases like real-time lakehouses and warehouses, no downtime migrations from NoSQL databases, real-time message streaming, and real-time cache ingestion.



## Key Capabilities

- Enables real-time data products
- Capture real-time streaming events and process pipelines
- Data capture from NoSQL data stores
- Continuous stream processing and transformation
- Optimized for microservices and analytics data stores
- Streaming visualizations and geolocation analytics
- Enables real-time distributed transactions
- Transform and shape data at scale

## Key Use Cases

- Transaction replication
- Data warehouse and data lakehouse ingest
- NoSQL data store migrations
- Real-time data transformation
- Oracle and non-Oracle SaaS integration
- Transaction Outbox pattern
- Big Data sources and targets

## Broad Topology Support

- On-prem to Oracle Cloud
- Oracle SaaS to Oracle Cloud
- Oracle SaaS to 3<sup>rd</sup> Party Cloud
- Oracle Cloud Cross Region
- On-prem to Autonomous Database
- Non-Oracle Cloud to Oracle Cloud
- On-prem to non-Oracle Cloud
- Oracle Cloud to Non-Oracle Cloud

**Data Lake connectors:** GG for DAA 23ai empowers data engineers and analysts to streamline real-time data delivery to data lakehouses. This eliminates data latency, enabling you to gain faster insights and make data-driven decisions and actions in real-time. GG for DAA 23ai seamlessly transforms your data into industry-standard formats like JSON, XML, Avro, ORC, and Parquet, ensuring seamless integration with various analytics platforms. This unique optimization for high-volume data ingestions ensures your data lakehouses are always up to date, empowering faster and more collaborative decision-making across your organization.

**Cloud Warehouse connectors:** GG for DAA 23ai supports various cloud warehouse technologies with stage and merge design which supports consistent and reliable data ingestions with very high throughputs. Cloud warehouse connectors provide no-downtime data migration capability with connected initial load and change data replication processes. These include OCI Autonomous Data Warehouse, Snowflake, Synapse, Big Query, Red Shift, Databricks, and more.

**NoSQL database connectors and high availability:** GG for DAA 23ai supports NoSQL technologies as sources and targets. Using GG for DAA 23ai you can configure MongoDB bi-directional replication for high availability use cases. NoSQL connectors include Oracle NoSQL, MongoDB (and derivatives: Atlas and Cosmos), Cassandra (and derivatives: DataStax), and more.

**Event and Messaging connectors:** GG for DAA 23ai supports Kafka (spanning Apache Kafka, OCI Streaming, Confluent Kafka/Cloud, AWS/Kinesis, Azure/Event Hub, Google Pub/Sub, etc.) and adds transaction integrity and reliability into this simple messaging technology. These connectors assure no-data loss, message order consistency and recoverability while replicating messages to streaming platforms. With this reliability, GG for DAA message streaming connectors become the backbone of mission critical applications and systems.

**Continuous Pipelines and Streaming Analytics:** GG for DAA 23ai also includes an Oracle Stream Analytics (OSA) license for data sourced by GoldenGate capture agents. OSA is complete platform for continuous integration pipelines (CTL) and real-time time-series and geospatial monitoring, and analytics. Pipelines support pattern matching, correlation, filtering, transforming, enriching, analyzing, alerting, and publishing data.

## Foundation for Reliable Distributed Microservices

Built for the demands of today's distributed data landscape, GG for DAA 23ai empowers developers to directly access and propagate real-time GoldenGate data events as well as support transaction consistency across modern microservices applications.

### Microservices reliability

Transaction Manager for Microservices Enterprise Edition (MicroTx) protects multi-datastore inconsistency when microservices distribute transaction management. Without automated coordination, it is a developer burden to ensure consistency across all the data stores. MicroTx removes the manual programming burden of ensuring consistency across all the data stores for all DML operations (updates, deletes, inserts) and failures.

### Key Benefits

- Single real-time data ingestion platform for microservices and analytics
- Extensive source and target connectivity across Oracle and non-Oracle technologies and all clouds
- Enable high-performance data replication with minimal impact to production systems
- Critical foundation for a real-time, heterogeneous, and distributed enterprise data fabric
- Enhance decision-making with trusted, real-time data
- Easy to use with enhanced user interfaces
- Access and integrate mission-critical applications without disruption

MicroTx is also useful when migrating from legacy platforms by keeping business processes in sync between legacy and modern apps during the migration process. In many cases, it is as easy as modifying a few lines of code modified or added to an existing microservice.

### Data Streams provides direct developer access to data events

GoldenGate Data Streams allows developers to access GoldenGate DML and DDL events using the AsyncAPI event framework within leading programming languages, such as Java, Python, .Net, PL, and JavaScript. Data Streams is more reliable than alternative messaging approaches, such as Kafka, due to GoldenGate’s fault-tolerance and exactly-once delivery.

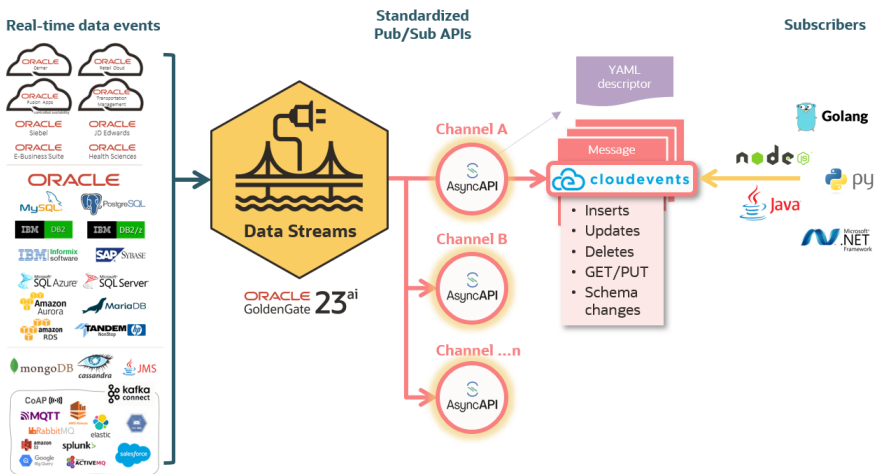


Figure 2: GoldenGate for Distributed Applications and Analytics AsyncAPI Support

## Conclusion

Traditional batch/bulk ETL data pipelines struggle to keep pace with the fast-paced world of microservices. GG for DAA 23ai emerges as a game-changer, specifically designed to bridge this gap. It empowers organizations to seamlessly integrate real-time data into their microservice architectures, fostering a event-driven data product culture for faster and more informed decision-making. This translates to a significant competitive advantage by enabling businesses to transform the customer experience landscape. GG for DAA 23ai encourages direct real-time event-based development with AsyncAPI support, ensures data consistency across distributed applications, and guarantees high-performance delivery without impacting source systems. By unlocking the true potential of real-time data within microservices, GG for DAA 23ai positions you for future innovation and success.

## Packaging and Supported Technologies

Please refer to [Oracle GoldenGate Licensing Information](#) for GG for DAA packaging details.

Please refer to [Configurations and source/target certifications](#) for source and target technologies supported by GG for DAA.

---

## Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com). Outside North America, find your local office at: [oracle.com/contact](https://www.oracle.com/contact).

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

---

Copyright © 2024, 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.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

---

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. 0120