

ORACLE

Oracle GoldenGate 23ai

Industry-leading database replication and data integration technologies
that connect and power your real-time data fabric

May 2024, Version 1.1
Copyright © 2024, Oracle and/or its affiliates
Public

ORACLE 23^{ai}
GoldenGate

Introduction

The Oracle GoldenGate (OGG) product family delivers real-time trusted information. For over 20 years, Oracle GoldenGate has been a leader for trusted Data Fabric / Data Mesh capabilities built around data product thinking, decentralized architectures, event-driven pipelines, and polyglot data payloads and topologies. Now, with Oracle GoldenGate 23ai, customers are able to leverage its transactional database replication for artificial intelligence use cases including Retrieval-Augmented Generation (RAG) and enabling real-time AI embedding.

The GoldenGate platform excels at change data capture (CDC), transactional data replication, data ingestion, data pipelines for continuous transformation and loading (CTL), and a wide variety of analytics on streaming data. GoldenGate interacts with the database transaction logs to continuously capture changed data (both DML and DDL) with minimal impact to the source database. This is why GoldenGate log-based replication is preferred over batch or periodically scheduled data captures. Only through leveraging such trusted data in real-time can organizations make better, faster business decisions and actions.

Oracle GoldenGate supports a real-time hybrid, heterogeneous data fabric that connects and unifies sources and targets via a multitude of styles for data integration processing:

- Real-time changed data capture (CDC) with transactional integrity
- Improve scalability and availability of mission critical databases
- High-speed ingest into big data targets
- Real-time continuous integration (CTL)
- Bulk ETL/ELT with Data Transforms in OCI GoldenGate
- Real-time streaming analytics for data-in-motion

Extensive Enterprise Connectivity

Oracle GoldenGate 23ai adds more than 30 new data sources and targets to allow connections to 100s of different databases, data stores, messaging services, applications, internet business services, and more. See the [GoldenGate and Data Transforms](#) documentation for the latest connectors.



Figure 1: Oracle GoldenGate supports 100s of connectors across clouds, databases, technologies, and application services.

Broad Portfolio of Solutions

The following Oracle products and cloud services comprise the Oracle GoldenGate product family:

- GoldenGate for Oracle databases
- GoldenGate for non-Oracle databases
- GoldenGate for Mainframe
- GoldenGate for Distributed Applications and Analytics
- OCI GoldenGate
- OCI Database Migration
- Stream Analytics
- GoldenGate Foundation Suite - Studio, Management Pack, and Veridata)

Key Capabilities

- Change data capture database replication
- Polyglot data integration with full heterogenous support
- Capture real-time streaming events and process pipelines
- Streaming visualizations and geolocation analytics
- Enables real-time Data Products and Data Mesh
- Continuous stream processing and transformation
- Transform and shape data at scale

Key Use Cases

- Replication of AI vectors and feed embedding algorithms with real-time data
- Transaction replication
- Data warehouse and data lakehouse ingest
- Database high availability
- Real-time data transformation
- Database migrations
- Event stream analysis

Real-time access to real-time information

Oracle GoldenGate software provides database high availability enabling business continuity for Oracle, MySQL, PostgreSQL, SQL Server, HPE NonStop, Db2 for i, LUW and z/OS. GoldenGate enables the highest service level, the Platinum Tier, of the Oracle Maximum Availability Architecture (MAA). Many thousands of global banks, retailers, telecoms, healthcare companies depend on Oracle GoldenGate to reliably run their operational data platforms.

The GoldenGate real-time ACID-based replication capabilities also offer an advantage for data movement for every purpose; across multivendor relational and document databases, across message stores, and across applications. Minimal source database overhead paired with high-speed processing make GoldenGate replication unique.

As a result, GoldenGate includes a large collection of certified heterogeneous sources and targets and is widely used for big data and NoSQL data transactions and analytical payloads as well as federating data into staging and search platforms.

Oracle GoldenGate 23ai is built as a microservices/REST based application. The modular architecture provides the flexibility for cooperative and distributed deployments across enterprise systems, on-premises and in the cloud.

What's New in Oracle GoldenGate 23ai?

Oracle continues to invest in all components of the Oracle GoldenGate product family, and key new features of Oracle GoldenGate 23ai include:

- Oracle Database 23ai support
- Vector database replication for Oracle, MySQL, PostgreSQL, Elasticsearch, and OpenSearch
- Embedded WebUI redesign optimized for large environments and fewer clicks
- New user experiences, extended develop Microservices with WebUI, RestAPI, and integration with 3rd party identity providers for authentication
- StatsD and Telegraf integration for real-time monitoring
- Data Streams via AsyncAPI with JSON output for CloudEvents
- New Oracle DDL commands to automatically add a table to replication
- More than 30 new heterogenous source/target certifications

Real-time data integration across the enterprise and cloud

Oracle GoldenGate 23ai captures and delivers real-time changed data to on-premises or cloud-based data warehouses, lakes and object stores, reporting systems, and other online transaction processing (OLTP) databases with minimal source system impact. Real-time data access improves business insight and analytics.

- **Oracle and non-Oracle data replication.** Oracle GoldenGate connects to many non-Oracle and open-source databases across on-premises and cloud platforms. Oracle GoldenGate also connects across multicloud relational database sources and targets such as databases in OCI, AWS Aurora and RDS environments, Azure

Broad Topology Support

- Integration with OCI GoldenGate as a source or target of data
- On-prem to Cloud
- Cloud to Cloud
- Cloud to on-premises

Key Benefits

- Enhance AI models and decision-making with trusted, real-time data
- Extensive connectivity across 100s of Oracle and non-Oracle technologies and services across on premises 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

SQL, and Google Cloud SQL and AlloyDB. Oracle GoldenGate supports many non-relational data platforms such as, Kafka, NoSQL databases, object storage, Hadoop, and more. Visit our [website](#) for a complete list of supported data platforms for each solution.

- **Real-time data warehouse.** To achieve continuous, real-time capture and delivery of changed data between operational and analytical systems.
- **Operational reporting and data integration.** Offload reporting activity from production databases to lower-cost secondary systems with current data for real-time reporting. Integrate operational data between OLTP systems in real-time.
- **Data Lakehouse ingest.** Provide continuous, real-time capture and delivery of changed data between OLTP and object storage, data lake, or data lakehouse targets. Oracle GoldenGate delivers data to Oracle or non-Oracle platforms, such as Snowflake, Azure Data Lake, Google BigQuery, and other cloud object stores.
- **Big Data integration.** Keep your entire distributed heterogeneous and multicloud data fabric and supporting applications up to date with real-time data replication, propagation, and transaction consistency. Simplify building and operating an event-based and zero data loss data fabric with five integrated enterprise-class components:
 - GoldenGate connectors to capture and apply data from NoSQL and messaging systems
 - GoldenGate Distributed Applications and Analytics connectors for object storage, data lake, data warehouse targets
 - Oracle Stream Analytics for continuous ETL and data in-motion analytics
 - GoldenGate Data Streams for developer control over all GoldenGate events
 - Transaction Manager for Microservices Enterprise Edition (MicroTx EE) for transactional integrity
- **Real-time streaming analytics.** Oracle Stream Analytics ingests real-time data from GoldenGate changed data, and from messaging environments like Kafka and JMS. A graphical pipeline designer can perform data transformations and apply time-series and geo-spatial analytics on these event streams. Oracle Stream Analytics is available standalone, a deployment option in OCI GoldenGate, and is included with Oracle Distributed Applications and Analytics.
- **Using GoldenGate in the cloud.** Oracle GoldenGate is available as a managed service in OCI, or as a self-managed OCI Marketplace image. Oracle GoldenGate on-premises can also be installed in any cloud on their compute instances.
- **Data Streams.** Most GoldenGate DDL and DML capabilities can also be accessed through developer APIs. Deployed using the AsyncAPI standard, developers can subscribe and publish to the GoldenGate data event model.

Maintain Continuous Availability of Critical Systems

Oracle GoldenGate helps organizations eliminate the downtime caused by both unplanned and planned outages while improving system performance and scalability. Oracle GoldenGate can be configured to support the following scenarios:

- **Zero-downtime operations.** Enable uninterrupted business operations during system, database, or application upgrade, migration, and maintenance activities.
- **Scalable Active-Active architecture.** Synchronize changes made across two or more databases to scale out workloads, provide increased resilience and near instantaneous failover across multiple geographic regions. Supports bidirectional replication with automated collision detection and resolution.
- **Data distribution.** Replicate data for distributed applications in real time across geographies for reliable access to timely data.
- **Query offloading.** Ensure high performance for production systems while supporting read-only activities by replicating data between heterogeneous sources and targets.

To learn more about GoldenGate's role in the Maximum Availability Architecture (MAA), read the best practices [documentation](#).

Conclusion

Oracle GoldenGate provides a fully managed real-time data fabric platform that helps organizations with real-time, continuous access to mission-critical information. Oracle GoldenGate addresses a wide array of continuous availability, disaster tolerance, data integration and streaming requirements. Available both on premises and in the Oracle Cloud, GoldenGate provides a modular foundation that scales to address the high-volume, low-impact data integration, streaming and replication challenges faced by enterprises today.

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

 facebook.com/oracle

 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