

Oracle GoldenGate Data Fabric

Capture once. Govern continuously. Activate everywhere.

October, 2026, Version 1.0
Copyright © 2026, Oracle and/or its affiliates
Public

Introduction

Applications, analytics, and AI systems depend on current, consistent business data. Yet change often crosses databases, cloud services, and data platforms on different schedules, leaving teams to reconcile movement, transformation, and governance across separate tools.

Oracle GoldenGate Data Fabric extends GoldenGate's committed-change foundation into a managed environment for building, governing, and operating data products. It connects change capture and replication with Data Pipelines, reusable Datasets, testing, access control, and operations—so teams can move from source change to products that applications, analysts, and AI consumers can use.




Figure 1: GoldenGate Data Fabric Overview

Delivered through Oracle Cloud Infrastructure (OCI) GoldenGate, the service is available in OCI and in supported multicloud environments such as Microsoft Azure and Google Cloud. Teams can choose processing modes, execution engines, storage locations, and data platforms to suit their architecture.

Key capabilities

- GoldenGate CDC integration with managed streams, extracts, and trails.
- Recipe-driven unidirectional, active-active, ZeroETL, and migration workflows.
- Streaming, micro-batch, and batch transformation pipelines.
- Streaming and tabular Datasets for processing and serving.
- AI-assisted design with test data, branches, proposals, and merges.
- Workspace boundaries with fine-grained role-based permissions.
- Data Topology and Replication Topology for operations.

Capture and transform change

Connect GoldenGate CDC to pipelines

Data Pipelines can use a GoldenGate CDC source as the start of processing. Data Fabric transparently generates and maintains GoldenGate artifacts—including data streams, extracts, and trails—so captured changes can flow into transformation pipelines.

Build replication workflows from recipes

Replication pipelines provide recipe-driven setup for common CDC workflows: unidirectional replication, bidirectional active-active, ZeroETL, and migration. Teams can choose a pattern that fits their source-and-target architecture, then oversee the resulting replication workflows as part of the broader GoldenGate estate.




Figure 2: Replication Pipeline Editor

Design pipelines for the target architecture

Data Pipelines process data in streaming, micro-batch, or batch mode. Users compose rich transformations such as joins, filters, and aggregations through AI prompts, low-code visual design, or high-code interfaces. Pipelines can mix storage locations and data platforms, with execution through Spark or SQL pushdown engines supported by the selected platform. Teams can use these choices to implement medallion layers, data lakes, lakehouses, and warehouses.




Figure 3: Data Pipeline Editor

Define Datasets for processing and serving

Datasets define reusable data interfaces for streaming or tabular data used in processing and serving. They can expose data from supported databases, Apache Iceberg catalogs, files, Kafka, and other stores. A Dataset connects its schema and endpoints with configured quality rules, lineage, processing evidence, and physical versions.

Teams can process data through pipelines and serve products to application, analytical, and AI consumers. Reusing a Dataset definition gives producers and consumers a common logical identity as the underlying data implementation evolves.




Figure 4: Dataset concept

Promote designs with a test-to-production workflow

Data Fabric supports a rich test-to-production methodology. Teams can generate test data, develop in branches, submit proposals, review and compare proposed changes, resolve conflicts, and merge accepted work before production. Previews and tests help evaluate pipeline changes and configured data-quality rules before promotion.

Bring AI assistance into design and operations

Models accessed externally can help users draft or refine pipeline designs through prompts, work through mapping and transformation decisions, and analyze production issues. Engineers can inspect and validate proposed changes before promotion.

Set project boundaries and role-based permissions

Workspaces group projects and separate their access controls while allowing approved resources to be shared. Role-based authorization provides fine-grained permissions for assets, controlling who can view or change them.

Operate distributed data architectures

Use topologies to understand and oversee data flows

Data Topology shows relationships among sources, pipelines, Datasets, and consumers so teams can understand how data moves through processing architectures. Replication Topology gives operators a way to monitor and oversee GoldenGate deployments and distributed replication workflows across the estate.

Choose deployment and execution options

OCI GoldenGate makes Data Fabric available in OCI and supported multicloud environments such as Microsoft Azure and Google Cloud. Teams can combine supported databases, Iceberg catalogs, files, Kafka, and other stores with Spark or SQL pushdown execution options. Exact combinations depend on supported sources, destinations, formats, catalogs, runtimes, and regions.

Representative use cases

- Operational replication and migration across supported systems.
- Active-active applications that need synchronized data copies.
- Streaming products for operational applications and services.
- Database change capture into Iceberg-based lakehouse layers.
- Medallion data processing and warehouse transformation.
- Governed data products for analytics and AI consumers.

Oracle GoldenGate Data Fabric brings replication, transformation, product definition, testing, and operations into a shared managed lifecycle. It gives teams a consistent path from captured business change to governed data products across supported cloud and data-platform choices.

Connect with us

To learn about GoldenGate Data Fabric, contact your Oracle account team or visit oracle.com.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2026, Oracle and/or its affiliates. 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. GoldenGate Data Fabric is in preview. Availability and supported sources, targets, runtimes, and catalog combinations vary by region and configuration.

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