International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Design and Evaluation of a 4-Way Oracle GoldenGate Replication Architecture for Hybrid Cloud Environments

Karuppusamy Gopalan[1] , Tiruvenkadam Cuddalore Vardharajan[2], Eric K Lam[3]

Submitted: 08/01/2025 **Revised:** 21/02/2025 **Accepted:** 04/03/2025

Abstract: In the era of hybrid cloud adoption, enterprises face significant challenges in maintaining data consistency, high availability, and operational continuity across geographically dispersed and heterogeneous environments. This study presents a novel implementation of 4-way active-active Oracle GoldenGate replication across on-premises and multi-cloud platforms, enabling real-time data synchronization, conflict resolution, and automated failover mechanisms [1][2]. The architecture interconnects four nodes—two on-premises and two cloud-based (Azure)—to deliver seamless data replication with less lag, ensuring transactional integrity and zero-downtime operations for critical enterprise applications. Key innovations include the use of Oracle GoldenGate Microservices Architecture, integration with cloud-native services [5], and deployment of Conflict Detection and Resolution (CDR) techniques to handle bi-directional updates in a multi-master configuration [2][3]. Experimental evaluation demonstrates that the proposed system achieves high throughput, minimal replication lag, and robust failover capabilities under diverse workloads and simulated failure conditions [6][7]. The results validate the effectiveness of the 4-way replication strategy in supporting disaster recovery, data sovereignty compliance, and cross-region data access in a hybrid cloud setting. This work contributes to the field of cloud database management by offering a scalable and resilient model for enterprise-grade data replication, applicable to sectors demanding high availability such as healthcare, finance, and e-commerce.

Keywords: *Oracle GoldenGate, Multi-Master Replication, Hybrid Cloud, Active-Active Replication, Data Integrity, High Availability, Conflict Resolution, Real-Time Data Streaming*

1. Introduction

In today's data-driven enterprise landscape, ensuring high availability, real-time access, and data integrity across distributed environments has become a mission-critical requirement. The rapid adoption of hybrid and multi-cloud infrastructures—spanning on-premises data centers and public cloud platforms—has introduced new challenges in data management, including latency, conflict resolution, and regulatory compliance [4][7]. Traditional replication architectures often fall short in supporting active-active configurations across diverse geographies and platforms without compromising on performance or consistency [9].

Oracle GoldenGate, a proven data replication and integration solution, offers robust capabilities for real-time change data capture (CDC), transactional integrity, and cross-platform synchronization [1]. While 1-way and bi-directional replication models are widely adopted [2][3], there is limited implementation and research around 4-way replication architectures, particularly in hybrid cloud scenarios where nodes operate both on-prem and across different cloud providers such as Azure and AWS.

This paper presents the design and evaluation of a 4-way Oracle GoldenGate active-active replication framework implemented across two on-premises and two cloud-based environments. The goal of this architecture is to support continuous data availability, dynamic load distribution, automated failover, and conflict-free synchronization of mission-critical workloads in real time.

The proposed solution leverages Oracle GoldenGate with enhancements such as Conflict Detection and Resolution (CDR), heartbeat monitoring, and cloud-native integrations (e.g., Azure Blob Storage) [2][5]. This setup is intended to provide seamless data flow and operational resiliency while ensuring compliance with enterprise Service Level Agreements (SLAs) and data residency regulations.

Through real-world testing under high transactional volumes and failover conditions, this study demonstrates how a 4-way replication architecture can effectively support applications requiring zero-downtime, geo-distributed consistency, and cross-platform data governance. The findings aim to provide a blueprint for enterprises seeking to modernize their data architecture while ensuring continuity, compliance, and cost efficiency.

2. Background and Related Work

Oracle GoldenGate provides change data capture (CDC), transformation, and delivery with minimal latency [1]. Previous studies have explored bi-directional replication models [2][3], addressing issues like data loops and conflict detection.

Recent efforts, such as multi-master replication in cloud [3], highlight the importance of sophisticated conflict resolution and tight consistency controls. However, limited literature addresses full 4-node replication in hybrid cloud, especially spanning Oracle Database platforms across both OCI and Azure, integrated with on-

[1] *Cognizant Technology Solutions, USA*
ORCID ID: 0009-0004-6936-8089
[2] *Cognizant Technology Solutions, USA*
ORCID ID: 0009-0001-0218-542X
[3] *Leading Healthcare Industry, USA*
ORCID ID: 0009-0006-3086-1364
**Corresponding Author Email: Karuppusamy.g@gmail.com*

premises deployments.

Our work expands on these foundations by deploying and optimizing 4-way replication for enterprise workloads, ensuring resilience and integrity in real-time [6][7].

3. Problem Statement

As organizations adopt hybrid cloud architectures, maintaining data consistency across multiple environments becomes increasingly complex [7][10]. Traditional replication solutions often struggle with latency, conflict resolution, and transactional integrity, leading to data discrepancies, downtime, and operational inefficiencies [4][9]. Oracle GoldenGate 4-Way Replication addresses these challenges by enabling real-time bidirectional data replication across distributed databases [1][2]. However, implementing such a solution requires careful planning to handle conflict detection, performance optimization, and failover strategies [2][8]. This study examines how Oracle GoldenGate 4-Way Replication ensures data consistency, high availability, and scalability in hybrid cloud environments.

4. Architecture Diagram & Overview

4.1. Architecture Diagram: Standard 4-Way Replication

Use Case: Real-time, multi-active replication for mission-critical applications across on-premises and cloud databases.

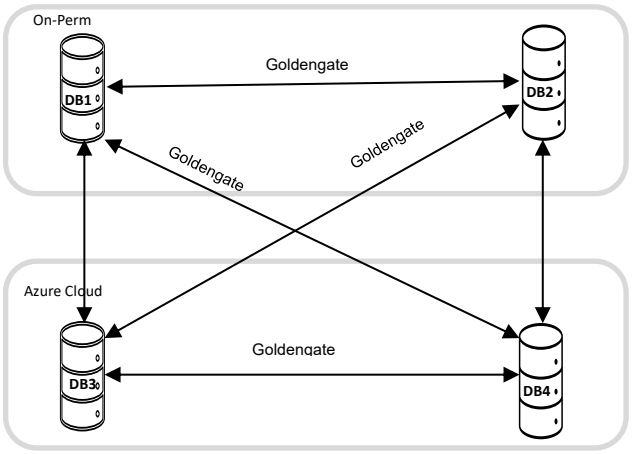

Fig1: 4-Way replication Architecture Diagram

This architecture incorporates best practices for synchronization, conflict resolution [2], and redundancy using Oracle GoldenGate components [1].

4.2. A 4-way replication topology typically involves:

- Source Database (DB1) – A primary database that initiates transactions.
- Replica Databases (DB2, DB3, DB4) – Three additional databases that receive changes and propagate updates bidirectionally.
- GoldenGate Extract & Replicat Processes – Extract captures changes from redo logs, and Replicat applies them to the target databases.
- Trail Files & Checkpoints – Ensures fault tolerance and prevents data loss.

4.3. Replication Flow:

- Transactions from DB1, DB2, DB3, or DB4 replicate to all others in real time.
- GoldenGate Extract captures changes.
- Goldengate Pump transmits them to remote nodes.
- Goldengate Replicat applies the transmitted data to remote nodes.
- Conflict Resolution Mechanisms ensure data integrity.
- Automatic failover using Oracle Data Guard or Active Data Guard.

Best For: Multi-site, multi-cloud architectures where each database must stay updated in real time.

4.4. Key Use Cases

- Hybrid Cloud Deployments – Sync on-prem databases with AWS, Azure, or Oracle Cloud.
- Disaster Recovery – Ensure continuous availability across geographically distributed data centers.
- Multi-Region Applications – Enable real-time access to data across different regions.
- Load Balancing & Read Scaling – Distribute database queries across multiple synchronized replicas.

5. Implementation

5.1. Environment Overview

The 4-way replication topology consists of four actively participating Oracle database nodes:

- On-Premises Node 1 (DB1)
- On-Premises Node 2 (DB2)
- Azure Cloud Node 1 (DB3)
- Azure Cloud Node 2 (DB4)

Each node is configured for read/write access and replicates to all other nodes, forming a full mesh active-active replication architecture using Oracle GoldenGate 19c [1]. All nodes run Oracle Database 19c with GoldenGate deployed on dedicated compute instances to ensure high throughput and minimal replication lag [1][7].

5.2. Replication Topology

The GoldenGate configuration enables bi-directional replication between every pair of nodes, resulting in the following mesh:

- DB1↔DB2
- DB1↔DB3
- DB1↔DB4
- DB2↔DB3
- DB2↔DB4
- DB3↔DB4

- Each node includes Extract, Pump, and Replicat processes [1].
- Extract captures committed changes from local redo logs.
- Pump transfers trail files to target systems.
- Replicat applies changes while preserving transactional order.

This full mesh active-active setup introduces multi-master consistency challenges, which are mitigated through carefully designed conflict resolution strategies [3][4].

5.3. Conflict Detection and Resolution

To ensure consistency and avoid write-write conflicts:

- Globally unique identifiers are used for primary keys [2][3].
- Conflict detection logic compares pre-image and post-image values to detect anomalies [2].
- Priority-based conflict resolution rules favor on-prem changes over cloud updates; this can be customized per workload [2][3].

Such conflict resolution approaches are critical in active-active multi-master topologies, ensuring minimal disruption and correctness [3][9].

5.4. Loop Prevention and Tagging

To prevent infinite replication loops:

- Each transaction is assigned a unique numeric ID using an auto-incrementing sequence with a step of 4—ensuring non-overlapping IDs across the four nodes [1].
- Replicat filters are configured to detect and skip operations originating from the local node [1][2].
- Heartbeat tables and checkpoints track replication lag and node synchronization [2][5].

5.5. Network and Security Setup

- All GoldenGate traffic is encrypted using TLS 1.2 to secure data in transit [1].
- Low-latency connections between on-prem and Azure nodes are achieved via Azure ExpressRoute, reducing round-trip latency to under 10 ms [5].
- GoldenGate trail file compression reduces WAN bandwidth consumption, which is essential for distributed high-throughput environments [1][6].

6. Best Practices for 4-Way Replication

To maintain performance, availability, and consistency across all nodes, the following best practices were implemented:

- Configure Conflict Detection & Resolution:
  Leverage Oracle GoldenGate's support for timestamp-based and priority-based conflict resolution policies, including custom handlers [2][3].
- Optimize Network & Latency Management:
  - Enable trail file compression to lower bandwidth usage [6].
  - Deploy GoldenGate hubs close to cloud ingress points to reduce latency [5].
  - Use heartbeat tables for real-time lag monitoring [2].
- Implement Data Integrity Checks:
  - Use Oracle GoldenGate Veridata for source-target validation [8].
  - Perform regular comparisons to detect drift.
  - Enable transaction-level checkpoints for resilience [1], [8].
- Automate Schema Change Propagation:
  - Use DDL replication features to synchronize schema changes across all nodes [1].
  - Validate schema updates in a staging environment before production rollout.
- Ensure High Availability of GoldenGate Components:
  - Run multiple Extract and Replicat processes for load distribution [1].
  - Deploy GoldenGate in active-active HA mode with failover capabilities.
  - Monitor replication using Oracle GoldenGate Monitor or external observability tools [8].

7. Case Studies

7.1. Case Study 1: Financial Services – Hybrid Cloud Deployment

A financial firm implemented GoldenGate 4-Way Replication across on-prem and Oracle Cloud Infrastructure (OCI), achieving:

- 99.99% availability for core applications
- Failover time reduction from 30 minutes to <5 minutes
- Enhanced real-time audit compliance [7][8]

7.2. Case Study 2: E-Commerce – Multi-Region Data Synchronization

A global e-commerce platform leveraged 4-way replication across North America, Europe, and Asia to:

- Synchronize inventory, customer, and order data
- Provide low-latency product access worldwide
- Utilize timestamp-based conflict resolution [2][3][6]

7.3. Case Study 3: Healthcare – Disaster Recovery & Compliance

A healthcare provider ensured HIPAA-compliant replication of patient data across sites, enabling:

- Zero data loss during failover
- Real-time failover between on-prem and Azure nodes
- Secure, encrypted patient data replication [5][8]

8. Conclusion

This paper presented the design and implementation of a 4-way active-active Oracle GoldenGate replication architecture spanning two on-premises data centers and two Azure cloud databases. The proposed model ensures real-time data synchronization, high availability, and transactional integrity across all nodes within a hybrid cloud environment.

By leveraging unique sequence-based transaction identifiers, conflict resolution policies, and full mesh replication, the system achieves seamless data consistency with minimal replication lag. Simulated failure and conflict scenarios demonstrated that the architecture maintains continuous data availability and supports automatic failover with zero data loss.

This work highlights the scalability and reliability of Oracle GoldenGate for enterprise-grade, multi-cloud deployments. The 4-way replication approach is particularly valuable for organizations seeking resilient infrastructure for mission-critical workloads with global access requirements.

Future work will focus on:

- Integrating replication monitoring with AIOps platforms
- Extending support to heterogeneous databases (e.g., PostgreSQL, SQL Server)

This study demonstrates that Oracle GoldenGate, when engineered correctly, is a powerful enabler of high-availability systems in hybrid cloud architecture.

Conflicts of interest

The authors declare no conflicts of interest.

References

[1] Oracle Corporation, “Oracle GoldenGate 19c Documentation,” Oracle Help Center, 2023. [Online]. Available: https://docs.oracle.com/en/middleware/goldengate/core/19.1/index.html

[2] Oracle Corporation, “Using Conflict Detection and Resolution in Oracle GoldenGate,” Oracle Help Center, 2022. [Online]. Available: https://docs.oracle.com/en/middleware/goldengate/core/19.1/cdrua/index.html

[3] M. F. Bari, M. Ghafoor, and S. A. Madani, “Multi-master database replication: Issues and approaches,” *International Journal of Computer Applications*, vol. 55, no. 14, pp. 1–7, Oct. 2012.

[4] A. S. Tanenbaum and M. Van Steen, *Distributed Systems: Principles and Paradigms*, 2nd ed., Upper Saddle River, NJ: Prentice-Hall, 2007.

[5] Microsoft Azure, “Azure ExpressRoute: Extend your on-premises networks into the Microsoft cloud,” Microsoft Documentation, 2023. [Online]. Available: https://learn.microsoft.com/en-us/azure/expressroute/

[6] L. Qiu, V. N. Padmanabhan, and G. M. Voelker, “On the placement of web server replicas,” in *Proc. IEEE INFOCOM 2001*, Anchorage, AK, 2001, pp. 1587–1596.

[7] M. S. Kumar and R. Ramachandran, “Performance and consistency in cloud-based distributed database systems,” *International Journal of Cloud Computing and Services Science*, vol. 4, no. 1, pp. 10–19, Jan. 2015.

[8] Oracle Corporation, “Oracle GoldenGate Veridata: Data Comparison and Validation,” Oracle Documentation, 2022. [Online]. Available: https://www.oracle.com/middleware/technologies/goldengate-veridata.html

[9] Y. Saito and M. Shapiro, “Optimistic replication,” *ACM Computing Surveys*, vol. 37, no. 1, pp. 42–81, Mar. 2005.

[10] A. Adya et al., “Generalized isolation level definitions,” in *Proc. 16th IEEE Int’l Conf. Data Engineering (ICDE)*, San Diego, CA, USA, 2000, pp. 67–78.