

# Performance Considerations for OCI GoldenGate

February 2024, Version [1.0]

Copyright © 2024, Oracle and/or its affiliates

Public

## **Purpose statement**

This document provides an overview of Oracle Cloud Infrastructure (OCI) GoldenGate features that can be used for data / delta lakehouse ingestion performance improvements.

It is intended to discuss different configurations for performance tuning, their potential impacts on your architecture, and guide you to the most optimal technical design for your business requirements.

## **Disclaimer**

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

## **Table of contents**

|  |          |
|--|----------|
| <b>Introduction</b>                    | <b>4</b> |
| <b>Replication performance factors</b> | <b>5</b> |
| <b>Performance tests</b>               | <b>6</b> |
| <b>Conclusion</b>                      | <b>7</b> |

## Introduction

OCI GoldenGate is the only cloud-native service providing the best-in-class data replication functionality with Oracle GoldenGate in a fully managed cloud environment. Using OCI GoldenGate, you can easily replicate your data securely, in real-time, at scale across heterogeneous cloud and on-premises systems.

OCI GoldenGate provides unique advantages:

- Auto scaling: When enabled, OCI GoldenGate deployment can automatically scale up to handle peak workloads.
- Upgrades and patching: With OCI GoldenGate, customers can apply patches and upgrades with a single click.
- Integration with native OCI Services: OCI GoldenGate integrates cloud-native OCI services (OCI Vault, OCI Monitoring, or OCI Networking) to provide unmatched cloud-native capabilities.
- Ease of use: OCI GoldenGate provides REST APIs for easy scripting and a comprehensive Web User Interface that allows for self-service configurations.

OCI GoldenGate supports dozens of different source and target technology combinations to support various use cases, from migrations to distributed analytics. Since OCI GoldenGate's release, real-time data ingestion into cloud storage services has been one of the most common use cases.

[OCI Object Storage](#), [Azure Data Lake Storage \(Azure Storage\)](#), [Amazon S3](#), and [Google Cloud Storage](#) are the cloud storage technologies supported by OCI GoldenGate. OCI GoldenGate supports JSON, delimited text, Parquet, Avro, and Avro OCF formats.

This technical paper focuses solely on performance considerations. Refer to the [OCI GoldenGate Documentation](#) and [Oracle Blogs](#) for other configuration details.



Figure 1. [OCI GoldenGate supported source and target technologies.](#)

## Replication performance factors

Three primary factors impact OCI GoldenGate replication performance. They are replicat type, network characteristics, and the GROUPTRANSOP parameter. This paper contains the results of twelve performance tests designed to help you optimize cloud storage ingestion performance.

Let's first discuss these performance factors.

### Replicat types

OCI GoldenGate provides two replicat types for cloud storage targets: *Classic Replicat* and *Coordinated Replicat*.

Classic Replicat is a single-threaded process that applies the messages to target cloud storage services.

Coordinated Replicat is a multi-threaded process where multiple threads read the OGG trail file independently and apply transactions in parallel. Each thread handles the filtering, mapping, conversion, and error handling for its assigned workload. By using [THREADRANGE](#) parameter in the replicat mapping, you can partition the workloads within the same replicat. Depending on the load in your tables, you can assign different threads for each mapping within the same replicat.

For example:

```
MAP test.table1, TARGET test.table1, THREADRANGE(1-3, ID);
```

```
MAP test.table2, TARGET test.table2, THREADRANGE(4-9, ID);
```

```
MAP test.table3, TARGET test.table3, THREADRANGE(9-10, ID);
```

### Network characteristics

Independent of OCI GoldenGate service performance, network characteristics including public and private connections, bandwidth, latency, and stability are consequential factors that impact replicat performance.

### GROUPTRANSOPS parameter

OCI GoldenGate cloud storage replicat process optimizes processing with *transaction grouping*. The [GROUPTRANSOPS](#) parameter groups multiple small transactions into a single larger transaction applied to cloud storage targets. The GROUPTRANSOPS parameter counts the database operations (inserts, updates, and deletes) and only commits the transaction group when the number of operations equals or exceeds the GROUPTRANSOPS configuration setting. GROUPTRANSOPS defers the transaction commit call until the larger transaction is completed. When a transaction is committed, the replicat flushes the operations.

The GROUPTRANSOPS parameter controls how frequently the flush process is called. By default, GROUPTRANSOPS is set to 1,000. Having a smaller value would result in more frequent flush calls, and having a bigger one would result in fewer. Fewer flush calls improve performance.

## Performance tests

Twelve test cases were run to understand how three different factors (GROUPTRANSOPS settings, replicat types, and network types) impact replication performance. The object storage target was an Azure Data Lake Storage using the Parquet format. Both the OCI GoldenGate source and Azure Storage target were co-located in the OCI Frankfurt data center.

The OCI GoldenGate deployment was configured with 4 OCPUs (the recommended minimum for production workloads), and auto scaling (which can increase the resources up to 3 times) was enabled. We used a single replicat process with the same data set from the same GoldenGate trail file in all our tests. Our data set represented OLTP workload and was made of 73,634,600 records, resulting in a 3.3 GB file in Azure Storage. OGG replication data transfer rates ranged from ~5 MB/sec up to ~40 MB/sec.

*Important note:* These tests were conducted to provide performance *guidance* for OCI GoldenGate and cloud storage services. Your replication performance **will differ** depending on the size and type of data replicated, the number of replicats you are running, cloud region locations, network settings, and replicat configuration settings. We recommend you work with an Oracle engineer and test your scenarios to find your most optimal configuration.

## Network Tests

Two network tests were conducted to evaluate the difference between a public and dedicated (OCI – Azure Interconnect 1 Gbps) connection. There was only a 3% difference between these two connections. Throughput for the dedicated network performed around 100 MB/sec and between ~5 MB/sec to ~40 MB/sec for the public network.

## Replicat Types

Classic and Coordinated Replicats were tested with two different GROUPSTRANSOPS settings. In general, performance improvements come at an additional memory cost. Coordinated Replicat uses more memory as it runs multiple threads. Each OCI GoldenGate OCPU resource has a fixed amount of memory, therefore additional OCPUs will meet also increase memory requirements.

Using *Coordinated Replicat* with 20 threads and *GROUPTRANSOPS* set to 1,000, there was nearly 5 times improvement compared to *Classic Replicat* with *GROUPTRANSOPS* set to 1,000. Coordinated Replicat used 8 OCPUs maximum for this test case.

Using *Coordinated Replicat* with 20 threads and *GROUPTRANSOPS* set to 20,000, there was nearly 3 times improvement compared to *Classic Replicat* with *GROUPTRANSOPS* set to 20,000. Coordinated Replicat used 12 OCPUs maximum for this test case.

## GROUPTRANSOPS Parameter

We tested both Classic and Coordinated Replicats with two different GROUPSTRANSOPS settings. In the first test, we used the default setting (1000). In the second test, we increased the GROUPTRANSOPS value to 20,000.

In *Classic Replicat* with *GROUPTRANSOPS* set to 20,000, we saw almost 3 times (300%) performance improvement compared to *Classic Replicat* with *GROUPTRANSOPS* set to 1,000. We saw only a marginal (~15%) change in resource utilization between these two cases. Both cases had 4 OCPUs allocated to GGS deployment.

In *Coordinated Replicat* with *GROUPTRANSOPS* set to 20,000, we saw nearly 50% performance improvement compared to *Coordinated Replicat* with *GROUPTRANSOPS* set to 1,000. In Coordinated Replicat, with *GROUPTRANSOPS* set to 20,000, the deployment had 12 OCPUs, while it used 8 OCPUs max with *GROUPTRANSOPS* set to 1,000.

## Conclusion

In our tests, increasing GROUPTRANSOPS setting provided better performance. For better performance/higher throughput consider using Coordinated Replicat. You can easily configure the GROUPTRANSOP and do your own tests. But also be aware that Coordinated Replicat will use more CPU resources because it increases number of replicat threads.

| Replicat Type   | Max OCPU | Performance Improvement |
|---|----------|-------------------------|
| Classic Replicat GROUPTRANSOPS 1,000                          | 4        | 1x                      |
| Classic Replicat GROUPTRANSOPS 20,000                         | 4        | 3x                      |
| Coordinated Replicat with 20 threads and GROUPTRANSOPS 1,000  | 8        | 5x                      |
| Coordinated Replicat with 20 threads and GROUPTRANSOPS 20,000 | 12       | 8x                      |

Table 1. The observed performance improvements for different replicat settings.

---

## Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **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

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail [REVREC\\_US@oracle.com](mailto:REVREC_US@oracle.com).

---