



Hewlett Packard
Enterprise

Performance test of real-time analytics with Oracle Database 12c In-Memory

Mixing OLTP and analytic workloads at large scale using the HPE Integrity MC990 X Server

Contents

Introduction.....	2
Test environment	2
Compute system	2
Storage.....	2
Oracle Database	2
Application.....	3
Data.....	3
In-Memory column store	3
Oracle Database In-Memory.....	3
HPE Integrity MC990 X Server	4
Performance tests	4
Mixed workload test results.....	4
Conclusion.....	6



Introduction

The innovative **dual-format** architecture of Oracle Database In-Memory enables you to accelerate analytic queries across your data warehouse by orders of magnitude. Non-standard queries of large tables that normally take several hours to complete, now take a matter of minutes. And there's much more. For example, Oracle Database In-Memory allows you to run transactional and analytic workloads on the same database concurrently. This helps eliminate the time and expense of ETL processes, and from a business line view, eliminate the delay between when analytic data is normally captured and what's taking place now. Imagine if your financial, manufacturing, and marketing teams could access vital information with up-to-the-moment accuracy, get answers instantaneously, and operate in real time.

Running a **mix** of transactional and analytic workloads on the same database can be game-changing—provided OLTP performance is not compromised. When leveraging Oracle Database In-Memory in a cluster, isolating analytic and OLTP workloads on different server nodes can therefore be considered to overcome performance overhead resulting from internode network latency. Alternatively, running both workloads on a single server can eliminate this latency. For large enterprises, a single server with many terabytes of In-Memory computing capacity is needed to run mixed workloads with multi-terabyte datasets.

Working jointly, SGI (recently acquired by Hewlett Packard Enterprise) and Oracle engineering teams conducted mixed workload testing of Oracle Database In-Memory using the Intel®-based HPE Integrity MC990 X Server (formerly named SGI UV 300RL) running Oracle Linux® and a database with over 10 terabytes and over one billion records. Intended for Oracle system architects, database experts, and IT personnel, this white paper presents the performance test results.

Test environment

The test environment for running transactional and analytic workloads concurrently using Oracle Database In-Memory consisted of the following:

Compute system

- HPE Integrity MC990 X Server with 32 sockets and 24 TB of shared memory running Oracle Linux 7.1 with the Unbreakable Enterprise Kernel (UEK) 3 (3.8.13-98.4.1.el7uek.x86_64)
- The HPE Integrity MC990 X Server contained 32 18-core 2.5 GHz Intel® Xeon® E7-8890 v3 processors with Hyper-Threading enabled and 768 32 GB ECC DDR4 DIMMs (24 per processor)

Storage

- SGI InfiniteStorage 5100
 - 60 x 900 GB SAS drives
 - RAID 6
- Connected to MC990 X using 16 x 8 Gb FC cables
- Oracle Automatic Storage Management (ASM)
 - 14 x 2.4 TB LUNs assigned to Oracle ASM (33.6 TB total)

Oracle Database

- Oracle Database 12c (12.2) with In-Memory option
- Total System Global Area (SGA)—16.1 TB
 - Fixed Size—9 MB
 - Variable Size—1.7 TB
 - Database Buffers—3.6 TB
 - Redo Buffers—1.1 GB
 - In-Memory Area—10.7 TB (for In-Memory column store and analytic/DSS processes)

Application

- Swingbench is a free load generator designed to stress test Oracle Database and provide performance data. The Swingbench Order Entry benchmark was utilized to create sales order data equivalent to a large enterprise and simulate high-volume sales transactions and analytic queries.
 - Order Entry is based on the **OE** schema that ships with Oracle Database 12c. It has been modified so that Spatial, interMedia schemas do not need to be installed. It can be run continuously, introduces heavy contention on a small number of tables, and is designed to stress interconnects and server memory.
- Swingbench is developed in Java. To learn more about Swingbench go to dominicgiles.com/swingbench.html.

Data

Swingbench was utilized to generate:

- An **ORDERS** table consisting of approximately 34 billion records totaling 3.71 TB and using 256 hashed partitions. The table consisted of 16 fields such as Order ID, Order Date, Customer ID, Address ID, Card ID, and Sales Rep ID.
- An **ORDER ITEMS** table consisting of approximately 103 billion records totaling 6.45 TB and using 256 hashed partitions. The table consisted of 11 fields such as Order ID, Product ID, Unit Price, Quantity, Gift Wrap, and Supplier ID.
- Additional smaller tables and table partitions such as **ADDRESSES, CUSTOMERS, WAREHOUSES,** and **LOG ON** totaling approximately 690 GB.
- Multiple indexes (21) totaling approximately 3.6 TBs to compare completion time for conventional row-based queries from disk with queries using the columnar format In-Memory.

In-Memory column store

Once the data was created, it was then populated (copied) in total into the In-Memory column store residing in DRAM (all tables were given the INMEMORY attribute with no column exclusions). The default level of In-Memory compression for Oracle Database In-Memory—MEMCOMPRESS FOR QUERY LOW—was selected as this optimizes for query performance vs. space saving. This resulted in approximately 3.6 TB in columnar format In-Memory vs. 10.8 TB on disk, or an average data compression ratio of 2.8.

Oracle Database In-Memory

With the introduction of Oracle Database In-Memory, a single Oracle Database can now efficiently support mixed workloads, delivering optimal performance for transactions while simultaneously supporting real-time analytics and reporting. This is possible due to the unique dual-format architecture that enables data to be maintained in both the existing row format for OLTP operations, and a new column format optimized for analytical processing.

Tables are stored on disk using Oracle Database's existing row-based formats. The new column format is a pure In-Memory format. The Oracle optimizer automatically routes OLTP operations to the row format and analytic queries to the column format, while allowing complete data consistency for all workloads. There is no need for persistent columnar storage (data already resides on disk in row format), so there are no additional storage costs or storage synchronization issues. The In-Memory column format also enables execution of ad hoc queries spanning billions of records in a matter of minutes. This can eliminate the need for most analytic indexes along with the performance cost associated with keeping indexes continuously updated during DML operations ((i.e., insert/update/delete).

Oracle Database In-Memory is easily deployed under any existing application that is compatible with Oracle Database 12c, and without application changes. To learn more about Oracle Database In-Memory, visit oracle.com/database/database-in-memory/index.html.

HPE Integrity MC990 X Server

Utilizing a unique scale-up, modular architecture and seventh-generation Hewlett Packard Enterprise technology, MC990 X is an advanced symmetric multiprocessing (SMP) system built for large and growing Oracle Database In-Memory environments. Featuring Intel Xeon E7-8800 v4 or v3 processors (up to 24 cores), the system is certified and factory-installed with Oracle Linux 7 with the Unbreakable Enterprise Kernel. A 5U modular chassis contains four sockets with up to 192 threads and integrated HPE NUMalink ASICs. By adding chassis (up to eight) and leveraging high-bandwidth, low-latency HPE NUMalink technology, HPE MC990 X can scale from four to 32 sockets and 1 to 24 TB of cache-coherent shared memory as a single system.

The single-system design of the MC990 X enables you to grow your Oracle Database In-Memory environment without adding overhead. There are no cluster nodes or cluster network to administer. No additional software is required. And there's no need for data segmentation as performance scales near linearly and automatically. HPE MC990 X can also be utilized to run Oracle Database 12c without the In-Memory option, as well as Oracle Database 11g. For more information about the HPE Integrity MC990 X server, visit hpe.com/servers/mc990x.

Performance tests

After generating the data set and tables outlined in the section titled **Test environment**, the Swingbench Order Entry (SOE) benchmark was utilized to first run OLTP and analytic workloads separately and then run the same OLTP and analytic workloads concurrently.

With Oracle Database In-Memory, OLTP workloads access data pulled from row format tables on disk into buffer cache like a conventional system. Analytic workloads run against data in columnar format in memory. These processes were run in parallel in the single MC990 X server when mixing workloads.

For OLTP, SOE simulated a high volume of sales orders over a five-minute period, which continually touches the large and small tables and partitions listed previously. The simulation was run three times, yielding minimal variation and the following baseline results:

- Transactions—49,788 total SQL statements
 - Selects—20,245
 - Inserts—13,325
 - Updates—10,432
 - Commits—5,786
- Average response—0.106 seconds

For analytics, SOE performed a **Sales Rep Query** of the 34 billion-record **ORDERS** table, which generated a report ranking the top ten orders for each rep by revenue. The query was also performed three times, also yielding minimal variation and with a baseline completion time of 1.8 minutes. While outside the focus of this white paper, In-Memory query times were consistently more than 50X faster than row-based queries from disk.

For mixed workload testing, the SOE sales order simulation and SOE Sales Rep Query were run in identical fashion, simultaneously. By comparing the respective OLTP and analytic workload results when mixed to running individually, an initial assessment can be made of the capabilities of Oracle Database In-Memory for real-time business in large enterprises. This is provided in the next section.

Mixed workload test results

Figure 1 illustrates a comparison between the number of sales order transactions completed when running separately (OLTP only) and when running concurrently with the Sales Rep Query (mixed workload). As shown, total OLTP transactions running over a five-minute period with a large-scale analytic query running concurrently were effectively the same—99.4 percent—as when running separately.

Figure 2 illustrates a comparison in response times between sales orders transactions running separately and when running concurrently with the Sales Rep Query. As anticipated based on the number of OLTP transactions completed, the analytic query had effectively zero impact on average OLTP response time.

OLTP SQL statements

(generated over 5 minutes)

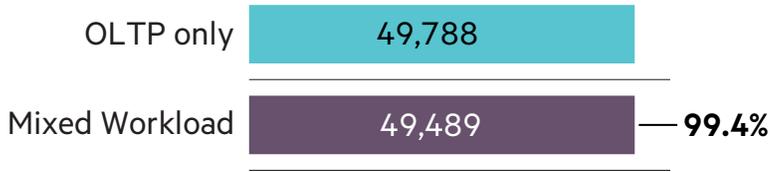


Figure 1.

OLTP average response time

(in seconds)

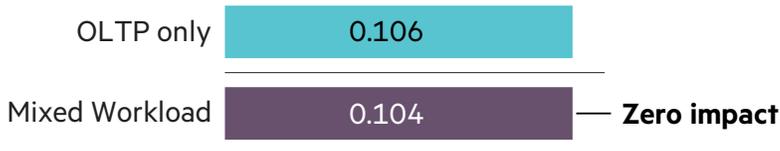


Figure 2.

The results in figures 1 and 2 collectively demonstrate the ability of Oracle Database In-Memory running on the HPE Integrity MC990 X to satisfy a key business requirement in this test environment—maintain transactional performance when running analytics in parallel. But what effect did mixed workloads have on analytic performance?

Figure 3 illustrates a comparison of the average completion times for the Sales Rep Query running separately (**Analytics only**) and running simultaneously with OLTP (**Mixed Workload**). As noted in figures 1 and 2, a full query of the 34 billion record **ORDERS** table took only 1.8 minutes to complete when running on its own. When the query was run simultaneously with OLTP transactions, the completion time did increase—but only by 24 seconds. Since OLTP transactions were unaffected by analytic operations (Figures 1 and 2), this comparison yielded two key findings with mixed workloads—large analytic queries continue to run orders of magnitude faster In-Memory vs. row-based queries from disk; performance degradation however slight was shouldered by analytics, not OLTP.

Analytics average completion time

(in minutes)

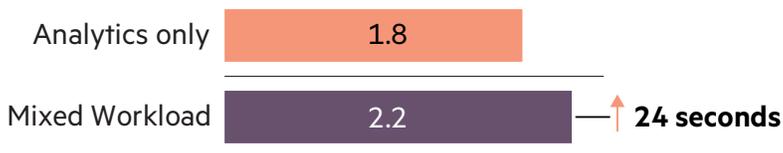


Figure 3.

Conclusion

This performance test conducted in October, 2016 provides a compelling indication of how Oracle Database In-Memory on a single HPE Integrity MC990 X server performs with mixed workloads running on Oracle Linux—large enterprises can run analytics on demand at extraordinary speed, without slowing transactions. Real-time analytics is now, and for business lines, it can mean significant competitive advantage.

Certainly, transaction volume and analytic completion times will vary as greatly as the workload and database attributes themselves. Accordingly, the data simulation, workload simulation, and test results outlined in this paper are not intended to substitute for actual customer testing. Hewlett Packard Enterprise therefore provides systems locally or remotely for customer POCs. Contact an [HPE Services representative](#) or your Hewlett Packard Enterprise partner to experience the breakthrough capabilities of Oracle Database In-Memory on the MC990 X with a test plan tailored to your database environment.

Learn more at

hpe.com/servers/mc990x



Sign up for updates

© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. All other third-party trademark(s) is/are property of their respective owner(s).

a00001520ENW, February 2017