

Oracle Advanced Compression Helps Global 500 Company meet Database Storage Savings Initiatives

Purpose statement

This document provides an overview of features and enhancements included in release 23ai. It is intended solely to help you assess the business benefits of upgrading to 23ai and planning for the implementation and upgrade of the product features described.

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

Executive Overview	4
Challenges and Expectations	4
Test Environment	4
Proof-of-Concept Testing, Results and Conclusions	5
Structured Data Compression (OLTP Table Compression)	6
Unstructured Data Compression (SecureFiles LOB Compression and Deduplication)	8
Backup Compression (Data Pump and RMAN)	9
Best Practices	13
Conclusion	13
Free Compression Advisor	13

About This Document

This Technical Brief provides an overview of the evaluation process a Global Fortune 500 corporation completed to test the various features included with Oracle Advanced Compression.

Executive Overview

Using their actual user data, Oracle database compression and space optimizations reduced the overall size of their disk consumption, on their proof-of-concept test systems, by 50% for structured data and up to 98% for unstructured data. As their testing demonstrates, the benefits from Advanced Compression can be significant, yet some organizations may still have uncertainty regarding the performance impact that compression may have on their systems.

This proof-of-concept confirms that Oracle Advanced Compression provides *storage savings and increased query performance* with no significant increase in overall CPU consumption. Following the completion of their proof-of-concept, the user purchased Advanced Compression and began to roll out their Advanced Compression deployment.

Oracle Advanced Compression typically provides database space savings as well as query performance improvements. In addition, given the growth experienced by today's databases, where operations like backups are increasingly taking longer to complete. Advanced Compression can help make those operational tasks faster, thereby providing both storage efficiency and query performance as benefits.

Challenges and Expectations

Exponential increases in data volumes in recent years have put enterprise IT infrastructures under severe pressure – from a cost, performance, scalability, and manageability perspective. It has become imperative to employ efficient ways of storing and managing data to meet the growing demands placed on IT systems.

The rapid growth in data volumes, and associated costs, were of significant concern – so much so that the user started global initiatives to identify and drive a reduction in storage costs.

Their IT organization identified data compression as an intelligent and efficient solution to address their data storage growth and costs challenges. As the user began to perform their due diligence, using Oracle Advanced Compression, their initial expectations began to take shape, which included:

- They expected their storage costs to drop as the result of any compression deployments, and
- They expected no query performance degradation and only minimal Data Manipulation Language (DML) performance impact from compression

Test Environment

The user performed proof-of-concept testing, and analysis, of Advanced Compression to determine if the capabilities provided matched their goals, namely, to minimize storage space usage/costs across Oracle database environments – with minimal or no query performance impact.

The test environment infrastructure included:

- Sun Sparc 64-bit w/ Solaris (96 GB RAM, 12 CPU)
- Oracle Database Enterprise Edition
- Quest Benchmarking Utility

Their proof-of-concept tested the following capabilities of Oracle Advanced Compression:

- Compression Advisor (DBMS_COMPRESSION)
- Structured data compression (OLTP Table Compression)
- Unstructured data compression (SecureFiles LOB Deduplication and Compression)
- Backup data compression (Data Pump export and RMAN Compression)

Proof-of-Concept Testing, Results and Conclusions

Compression Advisor Testing

Compression Advisor (available free with Oracle Database Enterprise Edition) provides an estimate of potential storage savings that can be realized using the Oracle Advanced Compression OLTP Table Compression feature. The estimate is based on analysis of a sample of data and provides a good estimate of the actual results that will be experienced when OLTP Table Compression is utilized.

The compression ratio achieved in each environment depends on the nature of the data being compressed, specifically the cardinality of the data. In general, Oracle Advanced Compression users can typically expect to reduce their storage space consumption by 50% or more by using Oracle Advanced Compression OLTP Table Compression feature.

The user tested Compression Advisor against approximately 20 to 30 tables with the goal of determining the expected compression ratios for the tables and then later comparing the estimates against the actual compression ratios, to determine the overall reliability of the compression advisor tool.

Below are the results of running Compression Advisor -- the test was performed on a table with about 60 million records taken from a production OLTP database.

Compression Advisor Output Example:

```
exec DBMS_COMP_ADVISOR.getratio ('TEST_USER','TEST_TABLE','OLTP',10);
Sampling table: TEST_USER.TEST_TABLE
Sampling percentage: 10%
Compression Type: OLTP
Compression ratio: 1.961 ←
```

The Compression Advisor output above shows, based upon a sampling of 10% of the TEST_TABLE data, the organization can expect a compression ratio of approximately 1.96x, or in other words, compression should reduce the storage requirements of the table by nearly 50%.

The actual compressed size of the table, after being built, proved to be a close match with advisor estimates, see below.

Compression Physical Testing (after table is compressed)

SEGMENT_NAME	MB	BLOCKS	COMPRESS	COMPRESS_FOR
TEST_TABLE	4405	563840	DISABLED	
TEST_TABLE_COMP	2268	290304	ENABLED	FOR ALL OPERATIONS

The compression estimate, produced by Compression Advisor based upon a 10% sampling, asserted a compression ratio of 1.96%, which means the 4405MB table would have compressed to 2237MB based upon

¹ The compression ratio achieved in each environment depends on the nature of the data being compressed, specifically the cardinality of the data.

the compression estimation. Once the table was physically compressed, the actual measured compression of the table saw the table compress from 4405MB to 2268MB, nearly identical to Compression Advisor's estimate.

The compression ratio achieved in this case indicates that the table has duplicate values that OLTP Table Compression was able to compress -- low cardinality data usually results in higher compression ratios. Further, the small difference between the Compression Advisor estimates and the actual OLTP Table compression enabled the organization to use Compression Advisor with a reasonable expectation that the estimates would be consistent with the actual compression as additional tables are compressed.

Structured Data Compression (OLTP Table Compression)

OLTP Table Compression often improves read operations, because database blocks stay compressed when read into the buffer cache. There is additional work performed while writing data, making it impossible to eliminate performance overhead for write operations. However, Oracle Advanced Compression attempts to minimize this overhead with the OLTP Table Compression feature.

With Oracle Advanced Compression's OLTP Table Compression feature, blocks are compressed in batch mode rather than compressing data every time a DML insert/update write operation takes place. A newly initialized block remains uncompressed until data in the block reaches an internally controlled threshold. When a transaction causes the data in the block to reach this threshold, all contents of the block are compressed in one operation. Subsequently, as more data is added to the block and the threshold is again reached, the entire block is recompressed to achieve the highest level of compression.

This process repeats until Oracle Advanced Compression determines that the block can no longer benefit from further compression. Only transactions that trigger the compression of the block will experience the re-compression overhead. Typically, DML OLTP transactions on compressed blocks should have similar, although some could be faster/slower, performance as before the block was compressed.

Test Scenario and Goals

No industry-specific test scenario was developed by the user for Oracle Advanced Compression structured data testing. Instead, the user utilized a pre-packaged, inventory management system test environment (open order, process order, delete order and then close order) that was included with their benchmark tool.

- Data was loaded connecting first to a user with the default tablespace uncompressed; after that test was finished the configuration was changed to a user with the default tablespace compressed, and the test was run again
- Testing consisted of loading 9 tables with indexes starting with 10 users and ending with 100

They had no pre-set expectations prior to beginning the proof-of-concept in terms of any overall performance impact because of compression -- but a primary goal of the proof-of-concept testing was to determine what, if any, performance impact would result from Advanced Compression.

To ensure that they tested for, and understood any performance implications on their OLTP system, the performance included not only QUERY testing, but DML UPDATE and DELETE testing as well.

Test Results and Conclusions: All times in minutes/seconds

Table Conditions	SELECT operation (select count (*))	DELETE Operation	UPDATE Operation (one column update)	UPDATE Operation (selected records update)
Not Compressed (60 Million Records)	28.04			12:06.36
Compressed (60 Million Records)	9.92			15.33
Not Compressed (100K)		5.93	4.51	
Compressed (100K)		7.56	2:13.38	

Source: User Tests

Conclusions

- Full table scan on “select count (*)” shows excellent results with compression. It took almost 1/3rd less time to do a record count of a 60 million record table (28.04 seconds vs. 9.92 seconds)
- DELETE DML on compressed data was ~20% slower than on uncompressed table (5.93 seconds vs. 7.56 seconds)
- One Column UPDATE DML showed slower performance with one column update for the whole table, while doing a selected records update shows 48 times better performance (last column in table) (4.51 seconds vs. 2:13.38 minutes).

Note: After their proof-of-concept testing Oracle helped identify some issues with bulk updates to OLTP table-compressed tables and these issues have been resolved with 11.2.0.2.

Note: There are some patches that can be applied to 11.2.0.1 or 11.1.0.7 that address these issues. See MOS Note #1061366.1 for the specific patches

- Selected Records UPDATE DML showed improved performance of 48 times better performance (12:06.36 minutes vs. 15.33 seconds)

Although not included on the table above, direct-path load tests showed no performance impact with Oracle Advanced Compression enabled.

Key Takeaways

Although the test scenario was deemed sufficient for the proof-of-concept, overall, they determined that the best test environment is where you can most closely duplicate the production environment with a test environment – but they acknowledge that due to time, budget, and infrastructure limitations this would have proved difficult to provide.

They felt that sufficient application testing is required to fully understand the performance characteristics of their data when using compression since performance cannot be 100% predicted.

Overall, they reported that most of their performance testing proved that the use of Oracle Advanced Compression's OLTP Table Compression feature had no, or very little, impact on their overall OLTP performance – in fact their test queries took 3x less time, deletes were not significantly impacted and selected record updates showed 48x improvement on compressed data.

The testing also helped them identify some areas, rarely encountered in their environment, as having a longer operations window with compression than without -- such as a one-column update across the entire 60 million record table.

Additional key takeaways realized, based upon their performance testing, include:

- Oracle Advanced Compression's OLTP Table Compression feature often improves read performance because database blocks stay compressed when they are read into the buffer cache (reduces memory use in the buffer cache)
- Redo generation could be up to ~30% more (based on the environment). This could result in more archive logs to be backed up, shipped to standbys, and increase recovery time with Oracle Advanced Compression
- Oracle Advanced Compression's OLTP Table Compression feature will have additional overhead for some DML operations. The overhead can vary based on the cardinality of the data, the CPU available and the Input/Output Operations per Second (IOPS) available

Unstructured Data Compression (SecureFiles LOB Compression and Deduplication)

Oracle Database's SecureFiles feature offers a 'best-of-both-worlds' architecture for storing unstructured content, such as documents, spreadsheets, and XML files. Oracle's SecureFiles feature is specifically engineered to deliver high performance for file data comparable to that of traditional file systems while retaining the advantages of Oracle Database.

Oracle Advanced Compression includes two storage optimization features for SecureFiles. SecureFiles LOB Deduplication is an intelligent technology that eliminates duplicate copies of SecureFiles data. SecureFiles LOB Compression utilizes compression algorithms to minimize the storage requirements of SecureFiles data.

Test Scenario and Goals

No industry-specific test scenario was developed by the user for Oracle Advanced Compression unstructured data performance testing. Instead, they utilized a small Java test (downloaded from Oracle Technology Network (OTN)) for read/write comparison between Oracle Advanced Compression's SecureFiles features for Deduplication and Compression versus BasicFiles.

BasicFiles is the Oracle feature, provided prior to SecureFiles, used to store large blocks of unstructured data. The tests involved comparing BasicFiles (no compression) storage to SecureFiles (with LOB Compression and LOB Deduplication)

Oracle's SecureFiles feature and LOB performance was measured using database driver protocol JDBC Thin driver (pure Java, Type IV driver)

Note that Test One and Test Two (see Test Results and Conclusions chart below) used the same SQL operations executed against two different data sets – one dataset provided by the Oracle OTN website and the other an internal dataset

Test Results and Conclusions: All times in seconds

Test Conditions	Compression and Deduplication	Used Space (bytes)	WRITE Operation	READ Operation
Test One				
BasicFiles	NO	4423680	1.17	0.13
SecureFiles	YES	49152	0.69	0.14
Test Two				
BasicFiles	NO	1048576	1.47	0.32
SecureFiles	YES	131072	2.14	0.26

Source: User Tests

Conclusions

- Space utilization utilizing Oracle Advanced Compression’s SecureFiles LOB Compression and Deduplication feature (vs. BasicFiles) was reduced by ~98%
- WRITE DML tests showed results of both faster and slower performance – additional testing would be needed to further define the specific performance characteristics of both
- READ operations show similar performance, or virtually no performance differences, between compressed and non-compressed tables

Key Takeaways

They had no pre-set expectations prior to beginning the proof-of-concept in terms of overall performance impact from SecureFiles LOB Compression/Deduplication.

A primary goal of the testing was to determine what, if any, performance impact would result from the use of Advanced Compression and to compare the effectiveness of compression/deduplication with SecureFiles LOB Compression/Deduplication to similar data stored in BasicFiles (without compression).

- Overall, they reported that most of their performance testing showed no, or very little, impact on their overall OLTP performance when using Oracle Advanced Compression’s SecureFiles LOB Compression and Deduplication features
- Additional analysis performed showed that they were able to achieve significant levels of deduplication – which not only surprised the proof-of-concept test team but also provided significant, and unexpected, storage savings
- Advanced Compression with SecureFiles LOB Compression and Deduplication is expected to enable them to reduce their unstructured data (large objects) storage requirements by approximately 98% (50x reduction) with no significant impact on performance – although compression did have some overhead associated with some WRITE operations

Backup Compression (Data Pump and RMAN)

With Advanced Compression, the capabilities of Data Pump compression were extended so that table data can be compressed on export with Advanced Compression. Oracle Advanced Compression’s Data Pump

compression feature is an inline operation, so the reduced dump file size typically results in a significant savings in disk space.

Unlike operating system or file system compression utilities, Oracle Advanced Compression’s Data Pump compression feature is fully inline on the import side as well, so there is no need to uncompress a dump file before importing it. The compressed dump file sets are automatically decompressed during import without any additional steps by the Database Administrator.

Oracle Advanced Compression also includes RMAN compression technology that can dramatically reduce the storage requirements for backup data. Due to RMAN’s tight integration with Oracle Database, backup data is compressed before it is written to disk or tape and does not need to be uncompressed before recovery – providing an enormous reduction in storage costs.

There are three levels of RMAN Compression: LOW, MEDIUM, and HIGH. The amount of storage savings increases from LOW to HIGH, while potentially consuming more CPU resources.

Test Results and Conclusions (Data Pump)

Data Pump Test Tables	Compression=ALL	Total Space to Export	Export dump size (GB)	Time (Min)	GZIP (GB)	Time to GZIP (Min)
Test One						
Uncompressed	NO	1.637	1.420	1.23	0.901	4.55
Uncompressed	YES	1.637	0.901	8.00		
Compressed	NO	1.518	1.420	1.22		
Compressed	YES	1.518	0.901	8.25		
Test Two						
Uncompressed	NO	4.301	3.808	1.40	0.831	5.20
Uncompressed	YES	4.301	0.672	9.40		
Compressed	NO	2.214	3.808	4.36		
Compressed	YES	2.214	0.672	14.54		

Source: User Tests

Conclusions

- No specific scenarios were tested as part of the backup compression testing. Instead, the proof-of-concept was intended to compare RMAN and Data Pump compression against the compression levels possible with 3rd party products like gzip

- Without Advanced Compression, Oracle Data Pump allows compression of metadata; with Advanced Compression, Data Pump also supports data compression in the export file
- Testing showed that Oracle Advanced Compression’s Data Pump compression feature compresses data slightly better than OS compression (GZIP) ~5% better
- They noted that Oracle Advanced Compression’s Data Pump compression feature compresses data as it does the export and does not need to have large space dedicated to uncompressed export files -- and that the import process understands compression in the export file
- The testing also showed that Oracle Advanced Compression’s Data Pump compression feature was slower than the OS (GZIP) compression (22-25%)

Note: By design, Oracle Advanced Compression’s Data Pump compression feature uses SQL to extract data from the database and export it, so it does not benefit from data already being compressed within the table.

If COMPRESSION=ALL is not listed as a parameter for export, the dump file will be created without compression.

Test Results and Conclusions (RMAN)

Tests were executed for RMAN Compression (included with RMAN) and RMAN with Advanced Compression (MEDIUM level) running backup and recovery on uncompressed and compressed tables.

They had no pre-set expectations regarding the performance of RMAN Compression versus RMAN with Advanced Compression; however, they believed it was important to monitor this test for any significant CPU overhead associated with backup compression.

Test results for previously uncompressed data

RMAN Test Conditions	Channels	DB Size (GB)	Backup Set Size (GB)	Compression Ratio	Time (Min)	DB GB/Hr
Test One						
RMAN Backup	8	105.9	18.04	5.9x	49:03	129.6
RMAN Restore	8	105.9	18.04	NA	43:42	
Test Two						
ACO Backup	8	105.9	24.6	4.3x	22:03	288.2
ACO Restore	8	105.9	24.6	NA	16:48	

Source: User Tests

Conclusions

- The test compared RMAN Basic compression (included with RMAN) and RMAN with Oracle Advanced Compression, running backup and recovery on uncompressed and compressed data – for Oracle Advanced Compression they tested using the MEDIUM level of RMAN compression

- Compression factor was 1.37x better using RMAN's Basic compression instead of RMAN w/ Oracle Advanced Compression
- Performance time was 2.22x better using RMAN with Oracle Advanced Compression instead of RMAN Basic Compression
- Database Restore/Recover Time was 2.63x better using RMAN with Oracle Advanced Compression instead of RMAN Basic Compression

Note: Test One utilized RMAN Basic Compression, Test Two utilized RMAN with Oracle Advanced Compression.

Test results for previously uncompressed data

RMAN Test Conditions	Channels	Compressed DB Size (GB)	Backupset Size (GB)	Compression Ratio	Time (Min)	DB GB/Hr
Test One						
RMAN Backup	8	74.3	13.04	5.7x	33.46	132.2
RMAN Restore	8	74.3	13.04	NA	32.10	
Test Two						
ACO Backup	8	74.3	19.7	3.8x	15.50	282.6
ACO Restore	8	74.3	19.7	NA	16:21	

Source: Customer Tests

Conclusions

- Compression factor was 1.50x better using RMAN Basic Compression instead of RMAN with Oracle Advanced Compression
- Performance Time was 2.26x better using RMAN with Oracle Advanced Compression instead of RMAN Basic Compression
- Database Restore/Recover Time was 1.98x better using RMAN with Oracle Advanced Compression instead of RMAN Basic Compression

Note: Test One utilized RMAN Basic Compression, Test Two utilized RMAN with Oracle Advanced Compression.

Key Takeaways

Overall, they reported that most of their performance testing showed that the use of Oracle Advanced Compression had no, or very little impact on their overall backup performance. Instead, testing showed that RMAN backups/recovery with Oracle Advanced Compression (MEDIUM level) was (on average) 1.4x faster for backups and 2.67x faster for recovery compared to RMAN backing up or recovering uncompressed datasets.

Specifically, their Proof-of-Concept testing demonstrated:

- With a compressed dataset the result backupset size with RMAN and Oracle Advanced Compression was almost the same as the backupset size of RMAN compression on uncompressed dataset

$$RMAN\ Compression\ (uncomp) - 18G / Oracle\ Advanced\ Compression(comp) - 19.7G$$

- With a compressed dataset the time to back up with RMAN and Oracle Advanced Compression was 1.42 times better than backing up the uncompressed dataset with RMAN Oracle Advanced Compression

Oracle Advanced Compression (uncomp) – 22:03m / Oracle Advanced Compression (comp) – 15:50m (initially ~105G DB size versus ~74G to backup)

- With a compressed dataset the time to recover with RMAN and Oracle Advanced Compression was 2.67 times better than recovering uncompressed dataset_with RMAN compression.

RMAN Compression – 43:42m / Oracle Advanced Compression – 16:21m

Best Practices

Below are some best practices identified in their Oracle Advanced Compression Proof-of-Concept:

- Each Oracle Advanced Compression feature requires thorough testing – the best test environment is where you can most closely duplicate the production environment with a test environment – but this is often difficult to do given budget, staffing and infrastructure limitations
- Space usage reduction with Oracle Advanced Compression enabled gives the best results where the most duplicate data is stored (low cardinality). This is especially true for backups -- greater compression will result in less data backed up and hence shorter recovery time
- Although CPU overhead is typically minimal, they determined that implementing Oracle Advanced Compression was preferred primarily on systems with available CPU cycles -- as compression will have additional overhead for some DML operations
- Tables can be redefined online, while queries continue against tables being converted to a compressed format

Conclusion

Following the completion of their Proof-of-Concept, they purchased Advanced Compression and rolled out their Oracle Advanced Compression deployments.

The first application to utilize Oracle Advanced Compression was an OLTP deployment for archiving/content management. They converted 40 database objects to utilize compression, with an average compression ratio of 2x (approximately a 50% reduction in space) -- which resulted in a reduction of the overall database size from approximately 8.1TB to 3.9TB.

They expect that the use of Oracle's Advanced Compression will play a significant role in helping them meet their storage savings initiative with no, or minimal performance impact.

Free Compression Advisor

An easy way to get started, with Advanced Compression, is by using compression advisor. The "DBMS_COMPRESSION" PL/SQL package (commonly called compression advisor) gathers compression-related information within a database environment. This includes estimating the compressibility of both uncompressed partitioned, and non-partitioned tables, and gathering row-level compression information on previously compressed tables/partitions. Compression advisor provides organizations with the storage reduction information needed to make compression-related usage decisions.

The output of running compression advisor is an estimation of the compression ratio for the specific table or partition that was the target of compression advisor. The output indicates the “COMPRESSION RATIO” presented as a number such as 2.1. This number indicates that, for this specific table or partition, the estimated compression ratio is 2.1x, which represents about a 50% reduction in the footprint of the table or partition should compression be enabled.

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. 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.

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