

Oracle Essbase Aggregate Storage Option Benchmark on Oracle's Sun SPARC Enterprise M5000 Server

*An Oracle White Paper
Jun 2010*

Oracle Essbase Aggregate Storage Option Benchmark on Oracle's Sun SPARC Enterprise M5000 Server

Executive Summary	3
Introduction.....	3
Benchmark Objectives	3
Product Overview.....	4
Benchmark Scenario Overview	5
Key Benchmark Metrics	6
Benchmark Details	6
Benchmark Infrastructure Profile.....	6
Storage Configuration	7
Test Application Characteristics and Configuration.....	8
Essbase Aggregate Storage Database Configuration.....	9
Test Methodology.....	10
Workload Specification	10
Benchmark Results Summary	12
Conclusion.....	19
APPENDIX.....	20

Oracle Essbase Aggregate Storage Option Benchmark on Oracle's Sun SPARC Enterprise M5000 Server

Oracle Essbase is the industry-leading OLAP (online analytical processing) server. Oracle Essbase provides an environment for rapidly developing custom analytic and enterprise performance management applications. With a rapid application development environment that is managed by the business, Oracle Essbase enables the business to quickly model complex business scenarios.

EXECUTIVE SUMMARY

Introduction

Business Intelligence (BI) systems are becoming more and more pervasive as stakeholders across organizations demand better transparency and visibility. To succeed in today's fast-paced, extremely competitive business environment, organizations have to constantly re-invent themselves. Success in today depends on an organization's ability to convert the vast amounts of data in various operational systems, data marts, and data warehouses into information and then derive intelligence from it. The key to addressing these business challenges is establishing a business intelligence infrastructure that delivers valuable information to employees, shareholders, customers and suppliers for timely, informed decision-making.

Oracle Essbase delivers key foundational technology to the business intelligence infrastructure for today's dynamic fast paced organizations. Oracle Essbase is the leading online analytical processing (OLAP) server available today. Oracle Essbase provides an environment tailored for business users to rapidly develop custom analytic and enterprise performance management (EPM) applications. Engineered for scalability, security, and rapid-response, Oracle Essbase brings analytics to the business user and enables the business to quickly model complex business scenarios. Through an intuitive interface, business users can manipulate large data sets to model complex scenarios, forecast outcomes, and perform what-if analyses to identify trends and optimize business results. Together with other modules of Oracle's EPM system, Oracle Essbase offers broad support for demanding business intelligence analytics requirements.

Benchmark Objectives

The objective of the benchmark is to showcase the ability for Oracle Essbase to scale in terms of user population and data volume for large enterprise deployments. Typical administrative and end-user operations for OLAP applications were simulated on Oracle's Sun Server and Storage products to produce benchmark results.

These benchmark test results demonstrate the following:

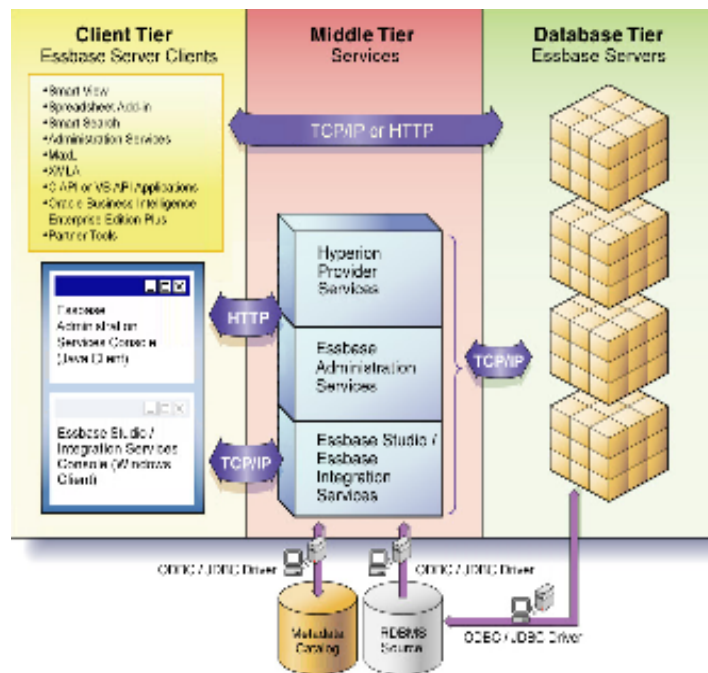
- Oracle Essbase can support a named user population of 200,000 users with the benchmark configuration. This assumes a 10 percent concurrency rate for the 20,000 active concurrent users in the benchmark.
- Oracle Essbase running as a 64-bit application deployed on SPARC64 VII powered Sun SPARC Enterprise M5000 server running Oracle Solaris 10 can handle a large scale enterprise sales analytics application with a database containing upwards of 1 billion records, 15 dimensions and dimensions with millions of members.
- Oracle Essbase can easily scale to meet the needs of any demanding business enterprise as demonstrated by the sub second query response time with 20,000 concurrent users and a large enterprise size Essbase application.

Product Overview

Oracle Essbase is a multi-threaded OLAP database server that takes advantage of symmetric multiprocessing hardware platforms. All Essbase application components, including database outlines and calculation scripts, application control, and multidimensional database information, reside on a server. Oracle Essbase provides two storage kernels for multidimensional databases, Aggregate Storage Option (ASO) and Block Storage Option (BSO). Aggregate Storage databases typically address read-only, “rack and stack” applications that have large dimensionality. Using Aggregate Storage, Essbase serves a wide range of analytic needs—financial analysis, planning, budgeting, sales analysis, marketing analysis, supply-chain analysis, and profitability analytics— all from a single analytic infrastructure. Block storage databases provide robust support for write-back and complex calculation operations and are well suited for modeling and forecasting applications. Essbase partitioning capabilities allow customers to take advantage of both storage technologies in a single application.

Oracle Essbase products incorporate powerful architectural features to handle a wide range of analytic applications across large multi-user environments. [Figure 1](#) provides a high-level view of the information flow between the three tiers of the Essbase architecture. The client tier includes Essbase Server clients, such as Oracle Hyperion Smart View for Office, Fusion Edition, Administration Services Console and Oracle Business Intelligence Enterprise Edition Plus. The middle tier includes services, such as Oracle Hyperion Provider Services and Oracle Essbase Administration Services. Communication between the client and middle tiers is through HTTP and communication between the client and database tiers can be either through TCP/IP or HTTP.

Figure 1 High-level Information Flow between Product Components



Benchmark Scenario Overview

The goal of this benchmark is to highlight how Oracle Essbase can support pervasive deployments in large enterprises on Oracle's Sun servers and storage, by simulating an organization that needs to support a large Essbase Aggregate Storage database with over one billion data items, large dimension with 14 million members and 20,000 active concurrent users, each operating in mixed mode: ad-hoc reporting and report viewing. The Oracle Essbase Aggregate Storage application employed for this benchmark was designed to model a scaled out version of a real customer financial business intelligence application.

The benchmark simulated typical administrative and end-user operations in an OLAP application environment. Administrative operations include: dimension build, data load, and database aggregation. Administrative operations were simulated using scripts through administrative interfaces available with Oracle Essbase. End-user testing modeled a total user population of 200,000 with 10 percent of the users actively retrieving data from Essbase. Scripts were used to simulate end-user reporting tasks. Think-time and time between iterations were built into end-user simulations.

The tests were performed on a single Sun SPARC Enterprise M5000 server running Oracle Solaris 10 Operating system and Oracle's Sun Storage F5100 Flash Array storage. These tests demonstrate scalability of Essbase and also provide a framework to enable customers to help reduce the risk, time, and complexity in deploying an OLAP solution.

Oracle Essbase supports extremely fast query response times for vast numbers of users, for large data sets, and for complex business models. The Oracle Essbase Aggregate Storage option demonstrates sub-second response times for a simulated environment with 20000 active concurrent users reporting against a financial business intelligence application.

Key Benchmark Metrics

This section summarizes the key metrics captured during this benchmark exercise on the Essbase Aggregate Storage database.

Test Name	Test Description	Test Results
Data Retrieval – 20000 Active concurrent users	In this test, user retrieval is simulated with MDX query scripts against a tuned Essbase database	Average query response time: <1 second
Dimension Build - Loading 13 million members to the “Customer” dimension	In this test, dimensions and members are loaded to an Essbase database outline by using a data source text files and rules file.	Total time to perform Dimension Build: 41 Minutes
Data Load into - 1 Billion Data Lines using 6 Load Buffers and 4 Data Load Threads.	In this test, data values are loaded to an Essbase database from a text files based data source and using rules file in parallel. Parallel data load was achieved with the help of multiple load buffers.	Overall Load Time: 269 Minutes
Database calculation through aggregation based on tracked queries	In this test, the data retrieval queries are tracked and resulting set of proposed aggregate views are materialized.	Elapsed time to perform the database aggregation based on tracked queries 115 Minutes

BENCHMARK DETAILS

The following sections describe the test infrastructure, test methodologies, and detailed results for the benchmarking exercise.

Benchmark Infrastructure Profile

This section outlines the infrastructure profile employed in the benchmark tests.

Component	Server specification
Oracle Essbase Server Oracle Essbase Administration Server	Oracle's Sun SPARC Enterprise M5000 Server <ul style="list-style-type: none"> • 4 X SPARC64 VII , 2.53 GHz (Quad Core) • 2 threads/core • Solaris 10 10/09 • 64 GB • 4x 1Gb Ethernet ports • Each core with own FPU • 4 X Local SAS drives with 10000 and 15000 RPM
Load Generation – LoadRunner Servers	Windows Client - 8 server s 8 CPU 24 GB Memory
Essbase Storage	Oracle's Sun Storage F5100 Flash Array Storage <ul style="list-style-type: none"> • Total 1 TB of disk storage – 40 Flash Modules • Sun SPARC Enterprise M5000 server attached to the F5100 via SAS HBA with 2 SAS channels • Oracle Solaris ZFS • Data Storage Scheme Striped – RAID 0 • 2 Oracle Solaris ZFS file systems created from a single zpool of 1 TB

The storage employed for this benchmark is the Sun Storage F5100, the world's fastest solid-state Flash Array storage. It provides unprecedented performance & eco-efficiency as well as delivers the most efficient way to accelerate database applications that doesn't require expensive tuning, power or space. It also decreases the latency and improves I/O response times for greater productivity and operational efficiency


The table below shows the product versions used in benchmark tests.

Product / Component Name	Version
Oracle Essbase	11.1.1.3 – 64 bit
Oracle Essbase Administration Services	11.1.1.3 - 64 bit
Oracle WebLogic	9.2 Maintenance Pack 64 bit JVM 1.5 update 19
HP LoadRunner ®	9.5

Storage Configuration

Sun Storage F5100 Flash Array Storage

Sun Storage F5100 Flash Array is attached directly to Sun SPARC M5000 server through SAS HBA. For this benchmark, 40 flash modules were employed with each flash module having a capacity of 22 GB and using these flash modules a single zpool was created. Two ZFS Filesystems were created from a single zpool.



Test Application Characteristics and Configuration

The Oracle Essbase Aggregate Storage application employed for this benchmark was based on a real customer application in financial industry with 13 Million members in the customer dimension.

This test application uses storage and dynamic hierarchy with up to 4-5 levels. All dimensions were structured as stored hierarchy with the exception of “measures” dimension which is structured as dynamic hierarchy and “Account” dimension is tagged as “compression” dimension.

Application details and outline:

```

- Outline: jumbo13m (Active Alias Table: Default)
+ Order_method Stored # Default # <4>
+ Order_date Stored # Default # <5> {Day_of_week}
+ Shipment_date Stored # Default # <5>
+ Shipper Stored # Default # <5>
+ Shipment_term Stored # Default # <8>
+ Sales_agent Stored # Default # <9>
+ Promotion Stored # Default # <12>
+ Pmt_term Stored # Default # <8>
+ Product Stored # Default # <4> {PSP_04, PSP_07, PSP_12}
+ Customer Stored # Default # <8>
+ Measures Accounts Dynamic Compression <8> (Label Only)
+ Day_of_week Attribute [Type: Text] # Default # <7>
+ PSP_04 Attribute [Type: Text] # Default # <25>
+ PSP_07 Attribute [Type: Text] # Default # <2>
+ PSP_12 Attribute [Type: Text] # Default # <12>

```

Dimension	Members in Dimension	Members Stored
Order_method	5	5
Order_date	1548	1546
Shipment_date	1548	1546
Shipper	6	6
Shipment_term	9	9
Sales_agent	811	810
Promotion	53	53
Pmt_term	9	9
Product	76379	76373
Customer	13284408	13284407
Measures	9	8
Day_of_week	8	8
PSP_04	26	26
PSP_07	3	3
PSP_12	13	13

The Oracle Essbase Aggregate Storage database was loaded with one billion lines of data from 40 different load files using different rule files to load into different intersections. The size of each load file was 1.2 Gigabytes containing 10 million data lines.

Essbase Aggregate Storage Database Configuration

This section outlines the different configurations applied to the Essbase environment in this benchmark exercise.

- Aggregate Storage Cache
 - Aggregate Storage cache facilitates memory usage during data loads, aggregations, and retrievals
 - Set to 32 GB – (half the total available memory on the Essbase Server)
- CALCPARALLEL set to 8, to enable parallel processing during view materialization.
- MAXLOGINS was set to 25000, to allow more than 20000 user sessions allowed to connect to the Essbase Server at any one time.
- Location of Tablespace and Source Files:
 - “Default” tablespace: “Default” tablespace was placed on ZFS Filesystem “DefaultTS” which use the zpool – “**ASODB**”
 - “Temp” tablespace: “Temp” tablespace was placed on ZFS Filesystem “TempTS” which use the zpool – “**ASODB**”
 - Source Files: Source files were placed on the M5000 internal drives (SAS drives with 10K rpm, striped with a stripe length of 128K).

Test Methodology

This section describes the test methodology employed to perform benchmark tests on the Essbase Aggregate Storage database. The Essbase Aggregate Storage database administrative functions were tested using custom developed MaxL scripts. Data sources for “Dimension Build” and “Data load” were flat files. The Data retrieval tests were performed using industry standard load test tool HP LoadRunner. In this test, proprietary test harness was used to submit queries directly to Essbase via the Essbase API.

Workload Specification

This section describes the workload employed in this benchmark testing exercise.

Aggregate Storage Database – Administrative Functions:

Dimension build:

This test involves, loading 13 Million members to the “customer” dimension using a flat files as a data source and rule file.

Data load:

Data was loaded into the Oracle Essbase Aggregate Storage database from 40 different load files using multiple rule files. The size of each of the load file was 1.2 Gigabytes containing 10 million data lines. This test was performed using a custom MaxL (Data Definition language for Essbase) script.

Database calculation:

Database is calculated by aggregating and storing the values.

Two options for aggregation were tested:

- **Default Aggregation:** Typically when retrieval request is submitted, Essbase calculates the data values by consolidating the values loaded into level 0 members these calculated values are not stored. In this test, the data values are pre-calculated and stored in the database by the aggregation process.
- **Usage Based Aggregation:** Fine tuning retrieval performance was performed with “usage based aggregation”. In this test, the proposed aggregate views which are created as a result of tracked retrieval queries were materialized.

Aggregate Storage Database – User functions – Data Retrieval:

This section outlines the query profile employed for the “Data Retrieval” tests. User requests were simulated using both Report Writer and MDX scripts.

Report Writer is a text-based script language that you can use to report on data in multidimensional databases. Following is the Report Writer request used in end-user testing.

```
{WIDTH 12} {OUTALT NAMES} <PAGE (Measures, Customer, Shipper)
QtyShipped <PARAM1> <Param2> <COLUMN (Shipment_Date) <CHILD
<Param 3> <ROW (Product) <CHILD <Param4> !
```

MDX is a language to develop scripts for applications to query and report against data and metadata in Essbase databases. Following are MDX queries used in testing.

Query 2: MDX Query 1:

```
SELECT {[<Param 1>].children, [<Param 1>]} on axis(0), {[psp_04].children}
on rows, TopCount
([customer].Levels(4).members,5,[Measures].[OrderLineDiscount]) ON
AXIS(2), {[Shipment_term].levels(1).members} on axis(3) from
<appname>.<dbname> where ([QtyShipped],[<Param 2>],[<Param 3>])
```

Query 3: MDX Query 2:

```
SELECT {[<param 1>].children, [param 1]} on axis(0), {[psp_04].children} on
rows, {[psp_07].children} on axis(2),TopCount
([customer].Levels(3).members,10,[Measures].[OrderLineDiscount]) ON
AXIS(3), {[Shipment_term].levels(1).members} on axis(4) from mill13m.mill13m
where ([QtyShipped],[<param 3>],[< param 4>])
```

Query 3 is executed once every ten iterations.

The different parameter values for the queries were supplied via data files and randomized during the test.

Data Retrieval test - step by step instructions:

1. User logs in to Essbase
2. User selects the application
3. Think time 1 second
4. User submits Report writer request
5. Think time 1 second
6. User submits MDX query 1
7. Think time 1 second
8. User submits MDX query 2
9. Logout

Steps 4 to 6 (both inclusive) were performed in a loop for the test duration of 1 hour.

Load Runner configuration:

- Random think time – 50 to 150%
- Pacing – Random - 60 to 90 seconds after each iteration

Benchmark Results Summary

This section describes the results of the benchmark tests performed with various configurations

Test 1 – Custom Dimension Load

In this test 13 million members were loaded into custom dimension in the Aggregate Storage database outline using rule files. This test was executed using a custom developed MaxL script.

Loading custom Dimensions involves building dimensions and performing a restructure. The elapsed time taken for these activities are show

Activity	Total Elapsed Time (in Minutes)
Building dimensions	37 Minutes
Restructure of the outline	4.2 Minutes

Test 2 – Data Load using Rules

In this test, 1 billion lines of data were loaded from 40 different load files using rules. Each load file was of size 1.2 Gigabytes containing 10 million data lines. For optimal performance, data values were loaded to a temporary data load buffer, and then written to storage after all data sources has been read.

The data load tasks involve the sub-tasks parsing of load file, loading the contents of load file to data load buffer and committing the contents of buffer to the database. In this test, the Sun Storage F5100 Flash Array Storage was employed.

The data load performance can be influenced by the following factors:

Number of Data Load Threads:

This configuration (DLTHREADSPREPARE) is typically set at the server level for all or specific Databases/application on the server. This setting enables parallel processing of data load files during the data load preparation stage which may improve data load performance. As a part of “Data Load” test, various executions were performed with different “DLTHREADSPREPARE” values.

Number of Load Buffers:

The data load performance can be improved by loading values to a temporary buffer load buffer, followed by a final write to storage after all the data load files are read. As a part of this “Data Load” test, various executions were performed with different number of data load buffers. The overall data load performance may be improved by having multiple load buffers and loading data into those buffers in parallel.

Aggregate Storage Cache:

Aggregate Storage Cache aids in data load and other Aggregate Storage Database tasks. In this benchmark test, cache 32 GB was used to determine the impact on performance.

Location of Data load files, Default and Temp Tablespace:

Tablespaces in Aggregate Storage Database help to optimize the data and work file storage and retrieval. Configuring the location of database (default tablespace) and work (temp tablespace) files on different physical drives helps to improve the performance. The data load files should also be placed on different drives to improve performance however it doesn't matter if it's on the same location as “default” tablespace.

Impact of Data Load Threads and Load Buffers on Data Load Performance

The impact of data load threads (DLTHREADSPREPARE) and the number of Load Buffers on data load performance were accessed by scheduling various test executions with different number of data load threads and Load buffers.

Storage Technology - location of tablespace & Data load file:

Sun Storage F5100 Flash Array Storage

- “Default” Tablespace: placed on the ZFS Filesystem “**DefaultTS**” which uses a single zpool “ASODB” of 1 TB.
- “Temp” Tablespace: placed on the ZFS Filesystem - “**TempTS**” which uses a single zpool “ASODB” of 1 TB.
- Data Load Files – Internal Drives of the Essbase Server running on the Sun SPARC Enterprise M5000 server

The metrics below demonstrates optimum data load performance can be achieved with 4 Data Load Threads and 6 Load Buffers. Please note: This optimum number of Data load threads and Load buffers depends on factors such as amount of physical memory and number of cores on the Essbase Server.

Number of Load Buffers and	Data Load – Elapsed Time (in minutes)
----------------------------	--

Data Load Threads	
# Load Buffers: 8 # Data Load Threads: 4	Overall Load Time: 282.3 Minutes Time to parse and write the contents of load file to buffer:128.77 Minutes Time to write the contents of buffer to Database: 153.60 Minutes
# Load Buffers: 6 # Data Load Threads: 4	Overall Load Time: 266.3 Minutes Time to parse and write the contents of load file to buffer : 141.72 Minutes Time to write the contents of buffer to Database: 124.65 Minutes

Resource Utilisation: Processor and IO Utilisation

The figures below show the processor and IO utilisation during data load test.

Execution 1: # Load Buffers: 6 # Data Load Threads: 4

Figure 1: Processor Utilisation










Metric	Average	Max	Min	Trend
CPU Utilization (USR)	44.00	94.00	0.00	
CPU Utilization (SYS)	1.00	5.00	0.00	
CPU Utilization (Total)	45.00	97.00	2.00	

Figure 2: IO Utilisation – Over all system activity (in seconds)

Metric	Average	Max	Min	Trend
Total Reads (per second)	79.34	2132.30	0.00	
Total Writes (per second)	139.23	891.60	0.60	
Total KB Read (per second)	10125.27	272933.50	0.00	
Total KB Write (per second)	15553.52	108030.70	15.60	
Average Service Time	20.44	153.50	0.00	
Average Disks used	30.00	44.00	3.00	

Execution 2: # Load Buffers: 8 # Data Load Threads: 4

Figure 3: Processor utilisation










Metric	Average	Max	Min	Trend
CPU Utilization (USR)	54.00	98.00	0.00	
CPU Utilization (SYS)	1.00	4.00	0.00	
CPU Utilization (Total)	55.00	100.00	0.00	

Figure 2: IO Utilisation – Over all system activity (in seconds)

Metric	Average	Max	Min	Trend
Total Reads (per second)	85.59	1474.20	0.00	
Total Writes (per second)	191.67	850.70	0.80	
Total KB Read (per second)	10898.52	188404.89	0.00	
Total KB Write (per second)	21794.79	104275.68	18.00	
Average Service Time	21.56	126.60	0.00	
Average Disks used	37.00	44.00	3.00	

Test 3 – Default Aggregation

In this test Aggregate Storage database is calculated by aggregating and storing values in the database. Aggregations are consolidations based on outline hierarchy and has two phases

- Aggregate view selection
- Calculation and storage of values for the selected aggregate views.

The factors influencing aggregation materialization performance are “Aggregate Storage Cache” and “Number of Threads” for parallel view materialization (CALCPARALLEL).

In this test for optimal performance, 8 threads for parallel calculation and 32 GB of Aggregate Storage cache were used. *Note: The number of threads for parallel calculation and amount of Aggregate Storage cache depends on the available hardware resources.*

Storage Configuration	Default Aggregation (Elapsed Time) in Minutes
Sun Storage F5100 Flash Array Storage Default” Tablespace placed on the ZFS Filesystem “DefaultTS” which uses a single zpool “ASODB” of 1 TB. “Temp” Tablespace placed on the ZFS Filesystem “TempTS” which uses a single zpool “ASODB” of 1 TB	526 Minutes (Total of 20 aggregate views built)

Resource Utilization: Processor and IO

The figures below show the processor utilisation of Essbase Server and IO.

Figure1: Processor Utilization

Metric	Average	Max	Min	Trend
CPU Utilization (USR)	14.00	26.00	0.00	
CPU Utilization (SYS)	0.00	3.00	0.00	
CPU Utilization (Total)	14.00	27.00	2.00	

Figure 2: This figure shows the over all IO activity

Metric	Average	Max	Min	Trend
Total Reads (per second)	81.32	1580.20	0.00	
Total Writes (per second)	6.11	57.70	1.20	
Total KB Read (per second)	10405.63	202246.39	0.00	
Total KB Write (per second)	69.96	3764.50	17.60	
Average Service Time	2.18	74.70	0.20	
Average Disks used	33.00	44.00	3.00	

Test 4 – Usage Based Aggregation

In this test, the aggregate views proposed as a result of tracked retrieval queries were materialized, to gain better performance for some retrieval requests.

Storage Configuration	Usage Based Aggregation (Elapsed Time) in minutes
Sun Storage F5100 Flash Array Storage Default” Tablespace placed on the ZFS Filesystem “DefaultTS” which uses a single zpool “ASODB” of 1 TB. “Temp” Tablespace placed on the ZFS Filesystem - “TempTS” which uses a single zpool “ASODB” of 1 TB	115 Minutes (Total of 6 aggregate views built)

Resource Utilization – Processor and IO

The sections below show the resource utilisation – Processor and IO. The figure below shows the processor utilisation of Essbase server during the test.

Figure 1: Essbase Server – Processor Utilisation

Metric	Average	Max	Min	Trend
CPU Utilization (USR)	14.00	26.00	0.00	
CPU Utilization (SYS)	0.00	3.00	0.00	
CPU Utilization (Total)	14.00	27.00	2.00	

Figure 2: This figure show the overall IO activity during the test.

Metric	Average	Max	Min	Trend
Total Reads (per second)	40.18	1180.90	0.00	
Total Writes (per second)	30.19	331.90	1.20	
Total KB Read (per second)	5135.77	150325.91	0.00	
Total KB Write (per second)	2368.31	38870.80	17.60	
Average Service Time	4.30	76.10	0.20	
Average Disks used	28.00	44.00	3.00	

Test 5 – Multiuser- Data Retrieval Tests

In this test, concurrent users submitted MDX (language-based data analysis mechanism to Essbase databases) queries to Essbase VIA API calls. In order to achieve concurrency industry standard load test tool HP® LoadRunner was employed. In this test scenario, test executions before and after usage based aggregation was performed.

Impact of usage based aggregation on retrieval performance

The metrics below clearly shows the performance improvements in data retrieval transactions after usage based aggregation. The queries submitted to Essbase were tracked using the “Query Tracking” option and the proposed aggregate views were materialized (usage based aggregation) for better retrieval performance.

Table 1 – Table below shows the transaction response time for Data Retrieval test executed before usage based aggregation with “Query Tracking” enabled. In this test, 3.38 queries per second and 3 reports per second were served during this 500 user retrieval test.

Transaction Name	Min (in seconds)	Average (in seconds)	Max (in seconds)	SD (in seconds)	90 Percent (in seconds)	Number of Transactions
Load Application	0.001	0.001	0.002	0	0.002	500
Load Database	0.001	0.001	0.008	0	0.001	500
Login to Essbase	0.004	0.006	0.027	0.001	0.011	500
MDX Query 1	0	3.64	65.953	10.306	4.239	25,620
Report Writer MDX Query 2	0	12.825	66.112	18.929	45.886	25,620
	0	61.132	184.563	41.599	113.625	2,421

Table 2 – Table below shows the transaction response time for Data Retrieval test executed after usage based aggregation. In this test, 4.9 queries per second and 4.43 reports per second were served during this 500 user retrieval test.

Transaction Name	Min (in seconds)	Average (in seconds)	Max (in seconds)	SD (in seconds)	90 Percent (in seconds)	Number of Transactions
Load Application	0.001	0.001	0.002	0	0.002	500
Load Database	0.001	0.001	0.002	0	0.002	500
Login	0.005	0.006	0.011	0	0.011	500
MDX Query 1	0	0.013	2.224	0.024	0.03	37,249
Report Writer MDX Query 2	0	0.003	1.83	0.021	0.01	37,249
MDX Query 2	0	0.082	0.918	0.092	0.19	3554

The rows marked bold are the response time for the queries submitted to Essbase.

Further data retrieval tests with 20000 concurrent users after usage based aggregation also showed similar performance with no signs of degradation in query timings. Table 3 shows the response time of the transactions with 20000 concurrent users.

Transaction Name	Min (in seconds)	Average (in seconds)	Max (in seconds)	SD	90 Percent	Number of transactions
Load Application	0.001	0.002	4.33	0.031	0.011	20,000
Load Database	0.001	0.004	4.967	0.095	0.011	20,000
Login	0.002	0.041	7.51	0.323	0.012	20,000
MDX Query 1	0	0.038	42.469	0.292	0.113	29,392,461
Report Writer	0	0.023	42.547	0.303	0.093	29,392,461
MDX Query 2	0	0.107	30.422	0.322	0.173	2,932,253

CONCLUSION

The Oracle Essbase Aggregate Storage benchmark demonstrates the ability of Oracle Essbase to support business intelligence applications with requirements for large data volumes and enterprise-scale user populations. The combination of Essbase, Sun SPARC Enterprise M5000 server and Sun Storage F5100 Flash Array Storage offer industry-leading performance for end-user and administrative operations in OLAP applications.

APPENDIX

Essbase - Aggregate Storage Database - Statistics

The section shows the Aggregate Storage statistics gathered at various stages of benchmark tests.

Essbase Database Statistics – after Data Load

Essbase Statistics	Value
Number of Input Cells	8041568592
Input level data size (GB)	54514176
Current Cache Size(GB)	32 GB
Disk space used by data (GB)	52 GB

Essbase Database Statistics – after Default Aggregation

Essbase Statistics	Value
Number of Input Cells	8041568592
Input level data size (GB)	54514176
Current Cache Size(GB)	32 GB
Disk space used by data (GB)	104 GB
Number of aggregate views	20
Number of aggregate cells	7146049384
Aggregate data size (KB)	53871424

Essbase Database Statistics – after Usage based Aggregation

Essbase Statistics	Value
Number of Input Cells	8041568592
Input level data size (KB)	54514176
Current Cache Size(GB)	32 GB
Disk space used by data (GB)	105.5 GB
Number of aggregate views	20
Number of aggregate cells	7146049384
Aggregate data size (KB)	56059968



Essbase ASO Benchmark
April 2010

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Copyright © 2010, Oracle. 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. Oracle, JD Edwards, PeopleSoft, Siebel, Hyperion, and Sun are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.