

Oracle Financial Services Revenue Management and Billing V2.3 Performance Stress Test on Exalogic X3-2 & Exadata X3-2

ORACLE WHITE PAPER | JANUARY 2015





Table of Contents

Disclaimer	1
Executive Summary	2
Introduction	3
Overview	3
Scope of Processing	4
Transaction Feed Management	4
Billing	4
Important Notes	5
Benchmarking Methodology	5
Setup and Configurations for Benchmarking	5
Data Loading	6
Processing	6
Measuring Tools	6
Results	8
Conclusion	10



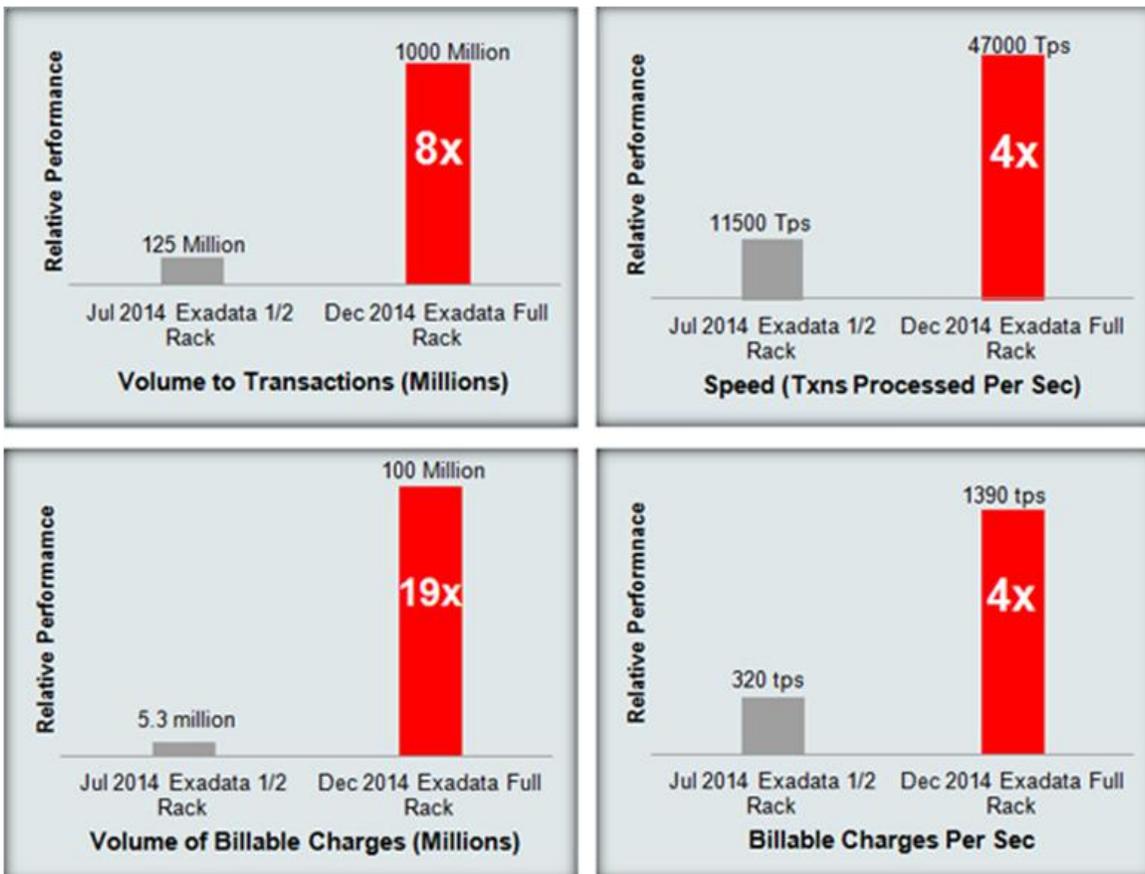
Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. 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, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Executive Summary

Oracle completed an extensive performance benchmark of Oracle Financial Services Revenue Management and Billing v.2.3 on Oracle's Exadata X3-2, for several key process workloads and customer/account volumes to evaluate linear scalability and throughput. The primary goal for this benchmark was to determine if the application can meet the performance requirements of the Banking, Payments and Insurance industries which lead to the processing of **one billion daily transactions and process 100 million billable charges for billing**. Apart from performance enhancements currently available in a service pack and planned for future release, Oracle Financial Services Revenue Management and Billing application was configured to run on an interval based partitioning and virtual column based sub partitioning strategy to achieve the performance goals.

The benchmark tested for an end to end process starting with transactions aggregation, pricing, producing billable charges and ending with producing bill segments and complete bills. When compared to a similar performance benchmark that was executed 6 months back, the new benchmark tested for 8 times the transaction volume and executed at 8 times the speed of the origination benchmark.



This document describes the performance and stress testing exercise conducted for Oracle Financial Services Revenue Management and Billing, our observations and conclusion for performance and scalability tests done for Transaction Feed Management(TFM), Billing and Pricing module executed on the Oracle Database Machine Exadata (full rack) and Exalogic servers (half rack) at Oracle India Development Center.

The following is a snapshot of the TFM and billing test runs. The same are detailed in the subsequent sections and testing and monitoring tool results have been included as applicable.

TRANSACTION FEED MANAGEMENT

Test Description	Avg TPS	Max TPS	Comments
1 Billion Transactions processed to generate five millions billable charges	44,500	52,000	Full rack of Exadata and 14 nodes of Exalogic were consumed

BILLING

Test Description	Avg TPS	Max TPS	Comments
8 million bills were created from 100 million billable charges	1,100	1,419	Full rack of Exadata and 14 nodes of Exalogic is used

Introduction

This white paper describes the performance and scalability results of Oracle Financial Services Revenue Management and Billing product version V2.3 on Exadata and Exalogic X3-2. From a performance perspective, Oracle Financial Services Revenue Management and Billing scaled to a billion plus transactions for usage of different products and pricing options for a range of few thousands to million customers, thus enabling clients to meet market windows and regulatory deadlines with low total cost of ownership and a high return on investment.

It demonstrates that should the occasion warrant the combination of Oracle Financial Services Revenue Management and Billing and Oracle Exadata can process more than **One Billion** transactions in less than six hours of processing window. The test also illustrates that the results scaled as expected when additional hardware was added and provides additional capacity that can address other needs.

Overview

Oracle Financial Services Revenue Management and Billing 's Transaction Feed Management enables the client to upload transactions received from various product processors or external applications and process them to determine the product usage of the accounts for billing based on the preconfigured Customer and Pricing setup.



The benchmarking exercise was performed on different configurations of Exalogic and Exadata to confirm the scalability of this performance - quarter, half and full rack. The Oracle Weblogic application server was used for deployment of application and Oracle Database as the database. Brief details about Oracle Financial Services Revenue Management and Billing, Exadata and application servers are covered below.

Oracle's flagship Billing product – Oracle Financial Services Revenue Management and Billing

- » This enterprise class, market-leading application offers advanced customer and financial data processing, revenue management, pricing & billing for the Banking, Insurance and Payment industries.
- » The test provides the application's rating to process complex transactions using business rules, products and pricing criteria to produce billable units with various sets of charges.
- » Processing and validating incoming raw transactions into the consumable units for billing as per the business rules.

Scope of Processing

Transaction Feed Management

Initial product Derivation: - System derives the accounts and the products with pricing parameters based on which the pricing is done. For each of the transactions the system saves the mapping of transaction to the derived Account-Product-Pricing parameters.

Grouping and Aggregation: - A group is created for a set of product parameters to determine product pricing and aggregate transactions to create billable charge. Execution of this batch creates unique groups along with ids for set of pricing/ aggregation parameters.

Pricing: - Pricing of the product will be searched across the hierarchy (customer hierarchy and pricelist hierarchies) to find effective Pricing for the account and final product that will be billed.

Exception handling: - If the effective pricing is not available or active contract is not available, or pricing is available with ignore switch as true, it marks the transaction as error or ignore else the price is derived and stamped.

Billable Charge creation: - Based on the service quantity identifiers (SQI) configured for the product in a division the batch aggregates the transactions and calculates the SQIs. It generates billable charge for the product and updates the SQIs.

Mark Complete: - All the transactions are marked complete if the billable charge is created successfully and service quantities are updated. In case the billable charge could not be created due to various reasons such as no active contract available, the transactions are moved to error state.

Billing

Pending Bill Generation: - The Pending Bill Generation generates pending bills for accounts that meet the criteria.

Bill Segment Generation: - The Bill Segment Generation generates the bill segments. Based on the batch business date, the system identifies the billable charges for which bill segments are to be generated. Once the billable charges are identified, the bill segments are generated in the pending bills.

Post Processing: - The Bill Completion generates post processing bill segments.

Bill Segment Freeze: - This process freezes the pending bill segments.

Complete Bills: - Once every Bill Segment of Bill is frozen, Bill is marked as complete. The financial entries are also posted during bill completion process.

Important Notes

The batch runs have been executed for a business day, which is an end of quarter or end of month billing cycle. The coverage includes complex processing scenarios. A typical business day is not expected to have such complicated use cases occurring on the same day and hence should take less time to execute as compared to the stress testing batch scenarios.

Benchmarking Methodology

The entire exercise comprised of system configuration for processing data, loading the transactions into staging table and processing the transactions to generate the billable charges. Methodology followed for each activity along with details of configuration/processing is covered below.

Setup and Configurations for Benchmarking

The system uses in-built rule engine to derive the accounts, product(s) and price variance parameter(s) associated with the incoming transaction based on the transaction attributes. Pricing setup included a price list with the product pricing assigned at account level for all the accounts. System performance was measured with multiple combination of following entities listed in the table.

ORACLE FINANCIAL SERVICES REVENUE MANAGEMENT AND BILLING ENTITY SETUP

Entity	Volumes
Number of Accounts	12 Million
Number of Customers	12 Million
Number of Divisions	31
Number of Products	600
Number of Rules Criteria	900
Number of Price Lists	5
Price Variance Parameters	10
Price Assignments	1.7 Billion
Number of Contracts	24 Million

Data Loading

Total of 300, 500 and 1,000 million transactions were loaded into the staging table. Transactions were distributed equally across 10,000 distinct accounts.

Processing

Each transaction was first validated for data integrity. For benchmarking purpose data was built up to ensure none of the transactions go into error state during validation to test the peak loads conditions. Validated transactions were mapped to the product and product parameters using the rule engine. Billable charges were generated for each set of distinct account, division, product and price variance combination. More than 5 million distinct billable charges were generated in the process after aggregation. Below is the time taken by TFM jobs to process the transactions (300 million to 1,000 million) and created 5 million billable charges.

Number of incoming raw transactions (Million)	Total Time (minutes)
300	104
500	169
1,000	321

Measuring Tools

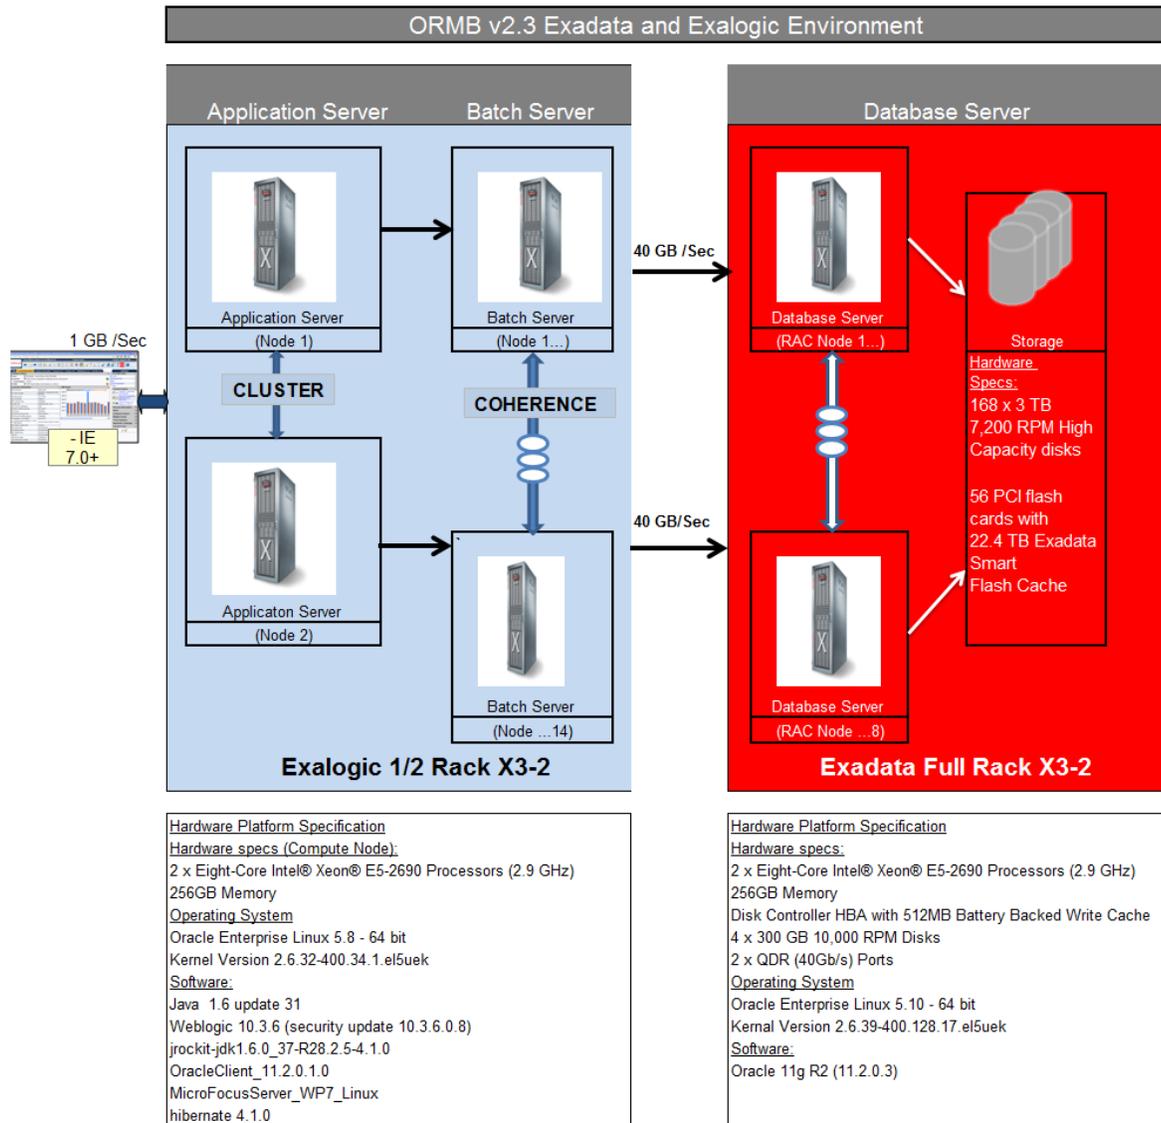
Oracle Enterprise Manager: A continuous monitoring of Oracle Enterprise Manager is performed to examine the state of the database and identify potential database performance bottlenecks even for fraction of a second. Any disparities observed were analyzed in detail and Oracle performance team took corrective action.

NMON: NMON is a tuner, benchmark and system administrating tool. NMON displays system performance information such as CPU, memory, hard drive and network traffic. Tool was run on every execution of batch. The output was collected in the form of spreadsheet.

AWR Report: The AWR (Automated Workload Repository) is used to collect performance statistics including wait events, Active Session History (ASH) statistics, object usage statistics and resource-intensive SQL statements. AWR snapshots are a source of information for several other Oracle features.

Topology & Configuration

This section shows the hardware topology of the systems used for the test as well as hardware and software combinations used.



THE COMPUTE NODES OF EXADATA/EXALOGIC SERVER WERE USED IN THE FOLLOWING MANNER:

Sr. No.	Purpose	Compute Nodes	Cores	Max Heap Size	Comments
1	Presentation Application Server	2	32	12 GB	A cluster of 2 Web logic Servers
2	Batch Server	14	224	1350 GB	A cluster of 14 nodes and 14 batch servers on each node.
3	Database Server	8	128	NA	RAC implementation of 8 nodes of Exadata.

Results

The test has shown linear scalability with the additional hardware. There were some changes to handle huge transaction in table like partitioning and sub-partitioning to prune the old transaction and mitigate the impact for subsequent runs.

BENCHMARK METRIC FOR TRANSACTION PROCESSING – JANUARY 2015

Details		Round 1**	Round 2	Round 3	
Hardware	Database Server	Machine	X2-2 Exadata 1/2 Rack	X3-2 Exadata 1/2 Rack	X3-2 Exadata full Rack
		Cores Used	4 nodes x 12 cores each	4 nodes x 16 cores each	8 nodes x 16 cores each
		RAM used	4 nodes x 96 GB each	4 nodes x 256 GB each	8 nodes x 256 GB each
	Application Server	Machine	Non EXA	X3-2 Exalogic-1/4 Rack	X3-2 Exalogic-1/2 Rack
		Cores Used	2 nodes x 32 nodes each	2 nodes x 16 cores each (< 1/4 rack)	14 nodes x 16 cores each
		RAM used	512 GB	2 nodes x256 GB (< 1/4 rack)	8 nodes x 256 GB each
Volume Processed (Incoming transactions)		128 million	125 million	1 billion	
Time Taken		6 hrs	3 hrs	6 hrs	
Speed (Transactions processed per second)		6000	11500	47000	
Scope of processing		receive txns >> product derivation >> Pricing >> Error Marking >> Billable Charge creation >> Mark Completion			
	Rules Criteria	325	325	900	
	Total Accounts	380,000	380,000	12 million	
	Price Points (Product x Account x Customer)	29 million	29 million	1.7 billion	
	Avg Price hierarchy	3 levels	3 levels	3 levels	
	AggTaper	125 million to 4.2 million	128 million to 5 million	1 billion to 5 million	
	Agg functions	2	2	2	

** At client location

TFM TPS BREAKUP BATCH WISE (1 BILLION TRANSACTIONS)

Test Descriptions	Avg TPS	Max TPS	Comments
Initial Product Derivation , Grouping and Aggregation is performed on one billion raw transactions	98,000	119,356	C1-TXNIP batch executed with 645 threads
Pricing is performed on all distinct accounts and product parameter combinations	22,000	25,983	C1-TXNVP batch executed with 100 threads
Mark Error	3,70,000	3,82,313	C1-TXNEX executed with 500 threads
Billable Charge creation	1,500	1,934	C1-TXNSQ executed with 150 threads
Mark Complete	167,500	189,387	C1-TXNCM executed with 500 threads

BENCHMARK METRIC FOR BILLING AND PRICING – JANUARY 2015

Details		Round 1**	Round 2	Round 3	
Hardware	Database Server	Machine	X2-2 Exadata 1/2 Rack	X3-2 Exadata 1/2 Rack	X3-2 Exadata full Rack
		Cores Used	4 nodes x 12 cores each	4 nodes x 16 cores each	8 nodes x 12 cores each
		RAM used	4 nodes x 96 GB each	4 nodes x 256 GB each	8 nodes x 256 GB each
	Application Server	Machine	Non EXA	X3-2 Exalogic-1/4 Rack	X3-2 Exalogic-1/2 Rack
		Cores Used	2 nodes x 32 nodes each	2 nodes x 16 cores each (< 1/4 rack)	14 nodes x 16 cores each
		RAM used	512 GB	2 nodes x256 GB (< 1/4 rack)	14 nodes x 256 GB each
Volume Processed (Billable charges)		4.3 million	5.3 million	100 million	
Time Taken		8 hours	4.5 hours	20 hours	
Speed (Billable charges processed per second)		150	320	1390	
Scope of processing		billable charge >> rating >> bill segment generation >> Bill completion & post processing >> Financial posting			
	Total Accounts	380,000	380,000	12 million	
	No of Bills	14,000	14,000	8 million	
	Contracts	800,000	800,000	25 million	
	Price Components	Flat, Step or Flat, Threshold	Flat, Step or Flat, Threshold	Flat	
	Calculations Line per Billable charge	2	2	5	

** At client location

BILLING TPS BREAKUP BATCH WISE (100 MILLION BILLABLE CHARGES)

Test Descriptions	Avg TPS	Max TPS	Comments
Pending Bill Generation	390	496	C1-PNDBL batch executed with 220 threads.
Bill Segment Generation	1,500	2,065	C1-BLGEN batch executed with 540 threads.
Post Processing	10,500	11,511	C1-BLPPR batch executed with 400 threads.
Bill Segment Freeze	28,000	29,989	C1-BLGEN batch executed with 400 threads.
Complete the Bills	3,500	4,373	C1-BLPPR batch executed with 400 threads.

Conclusion

This benchmark demonstrated that Oracle Financial Services Revenue Management and Billing v2.3 is able to process transactions at a rate of 44,500 per second and process 100 million billable charges using compute nodes of Exalogic X3-2 as the application tier and a full rack of Exadata x3-2 server on the database tier.

Oracle Financial Services Revenue Management and Billing v2.3 application's ability to scale near-linearly was proven both vertically and horizontally. The TFM and billing job shows near-linear scalability up to the point where throughput becomes limited by hardware resources.

The results from this Exadata X3-2 benchmark shows an improvement of eight times higher processing rate relative to the prior baseline benchmark conducted. This improvement is determined to be primarily associated with the architectural & design changes made in the software for better throughput and memory utilizations. Few of the prominent contributors for this extreme performance are:

- » Model way of coupling the code with partitioning and sub-partitioning strategies using virtual columns for distributed processing across partitions with isolation and data pruning to achieve negligible concurrency, contentions and waits on database
- » Introduction of various caching (in memory) algorithms at various places for entity and transactional cache thus reducing the fetch time from database
- » Software optimizations to reduce the database I/O
- » De-normalization of pricing data model to avoid crawling in the complex pricing hierarchy tree resulted speedier price search
- » Minimization of analytical functions with tuning improved the execution speed for queries
- » Software tuning and parameterization to suit the hardware resources to ensure optimum performance with scalability
- » Leveraging the various performance features of Exadata hardware

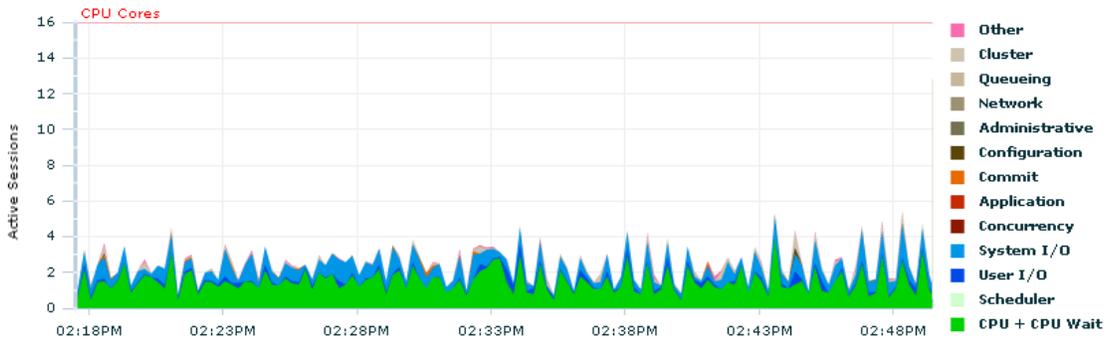
CPU AND MEMORY UTILIZATION

	Batch Server	Database Server
Max CPU Utilization	89 %	27%
Avg. CPU Utilization	77%	22 %
IOPS	-	15,410
RAM	1050 GB	736 GB
Heap Memory Utilized	54 GB	-
Throughput	-	620 MB/Sec

AWR Report - Foreground Wait Classes - % of Total DB time

CLUSTER-WIDE TOTALS OF WAIT CLASS FOREGROUND WAIT TIMES AS A PERCENTAGE OF THE CLUSTER-WIDE DB TIME

User I/O	Sys I/O	Other	Applic	Commit	Network	Concurrency	Config	Cluster	DB CPU
7.87	0.00	0.59	0.00	0.23	0.00	0.07	1.67	7.11	87.01



Observations:

- » Optimum DB CPU utilization
- » No Contention
- » Low cluster wait
- » No network wait

AWR REPORT - SYSTEM STATISTICS

Statistics	Total	Per Second	Per Transation
physical read IO requests	1,431,403	1,589.01	15.58
physical read bytes	55,955,931,136	62,116,826.98	608,957.98
physical read requests optimized	1,282,834	1,424.08	13.96
physical read total IO requests	1,452,830	1,612.79	15.81
physical read total bytes	56,302,891,520	62,501,988.60	612,733.89
physical read total bytes optimized	10,782,531,584	11,969,729.22	117,344.28
physical read total multi block requests	46,634	51.77	0.51
physical reads	6,830,558	7,582.62	74.34
physical reads cache	6,045,659	6,711.30	65.79
physical reads cache prefetch	4,627,076	5,136.53	50.36
physical reads direct	784,899	871.32	8.54
physical reads direct (lob)	1,323	1.47	0.01
physical write IO requests	11,845,441	13,149.52	128.91
physical write bytes	209,098,858,496	232,120,947.19	2,275,583.96
physical write total IO requests	12,428,827	13,797.14	135.26
physical write total bytes	530,061,833,728	588,422,660.65	5,768,564.27
physical read IO requests	1,431,403	1,589.01	15.58
physical read bytes	55,955,931,136	62,116,826.98	608,957.98

Observations:

- Throughput – 620 MB/sec
- IOPS – 15410



CONNECT WITH US

-  blogs.oracle.com/financialservices
-  facebook.com/oraclefs
-  twitter.com/oraclefs
-  oracle.com/financialservices

Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

Hardware and Software, Engineered to Work Together

Copyright © 2015, 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.

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.0115

White Paper Title: Oracle Financial Services Revenue Management and Billing V2.3 Performance Stress Test on Exalogic X3-2 & Exadata X3-2
January 2015
Author: Oracle Financial Services Revenue Management Product Marketing Team

