



An Oracle White Paper  
July 2010

# Oracle Utilities Meter Data Management 1.5 Performance Benchmark on Exadata 1

Introduction .....	1
Overview .....	2
Methodology.....	2
Historic Data.....	2
Load Data.....	2
Validation and Estimation .....	3
Workload Profile .....	3
Performance and Benchmarking Environments.....	3
Exadata Configuration .....	4
Application Servers.....	5
Result .....	6
Conclusion .....	10

## Introduction

Oracle Utilities Meter Data Management is an enterprise-class application suite built on open standards designed for high levels of automation and scalability for the transformation, storage, validation, estimation, and retrieval of meter data. A few key features of Oracle Utilities Meter Data Management that are relevant to this benchmark are the following:

- Provides a scalable solution to meter data load including smart meters with increasing capabilities to capture and communicate data.
- Provides configurable meter data validation and loading capability.
- Provides configurable estimation process for missing or invalid meter data.

As the number of smart meters in a utility's install base increase, the volume of data needing to be processed will also increase significantly. In the current market, utility companies are looking for a software and hardware solution that can provide them with scalability and flexibility that allows the business to grow smoothly as the amount of information and processing needs grow.

Smart meter deployments place a rapidly increasing demand on the performance, fault tolerance, and scalability of core utility company business applications. Oracle Exadata provides a unique solution to these specific challenges.

This white paper describes the performance and scalability results of Oracle Utilities Meter Data Management version 1.5.0.6 on the Exadata 2 full software stack using Exadata 1 hardware. It demonstrates the near linear application scalability up to the maximum CPU capacity of Exadata 1 database machine.

## Overview

Oracle Utilities Meter Data Management version 1.5 is a comprehensive product that provides transformation, storage, validation, estimation, and retrieval of meter reading data. It provides many out of the box validations and estimation routines as well as the capability to create custom rules as business needs dictate.

In this benchmark, Oracle Utilities Meter Data Management version 1.5 demonstrated the ability to run on the Exadata with near linear application scalability up to the max CPU capacity of a half-rack of Exadata 1 database server nodes.

## Methodology

### Historic Data

For the benchmarking effort, 13 months of historic data for five million interval meters was used. The total number of historic interval readings was 65 million reads with each read representing one month of 15-minute interval data. For each 24 hour period within the read there were 96 individual intervals. This equates to approximately 2880 intervals per each monthly read. This data was randomly generated using an internal data generation tool. This was necessary to ensure that data would not be cached by the database and inflate the throughput numbers. This approach is critical when testing with Oracle Secure Files and using table compression.

In addition to the interval historic data, another five million consumption meters and the associated 13 months of historic data for these consumption meters was also loaded into the database. The following table represents total number of rows per key entity in the database:

ENTITY TYPE	NUMBER
Account	10,000,000
Interval Meter	5,000,000
Scalar Meter	5,000,000
Total Meter	10,000,000
Service Point	10,000,000
Historic Interval Meter Reads	65,000,000

### Load Data

The test data load corresponds to 5 million daily cuts where one cut equals one day's data for one meter, 15 minute intervals. Thus a load includes total 480 million individual usage intervals. The test

data was comprised of XML files based on the LSM data format. Each file consisted of 1000 daily cuts for 1000 meters. The approximate file size for an interval test file was 3mb.

## Validation and Estimation

Additionally several out of the box validations were benchmarked with a pre-configured percentage of failures for each validation. The errors were randomly distributed across the entire 5,000,000 read data set. The following table lists the validations used in the benchmark as well as the percentage of failure for each validation. Again, the goal was to create as close to a real world scenario as possible.

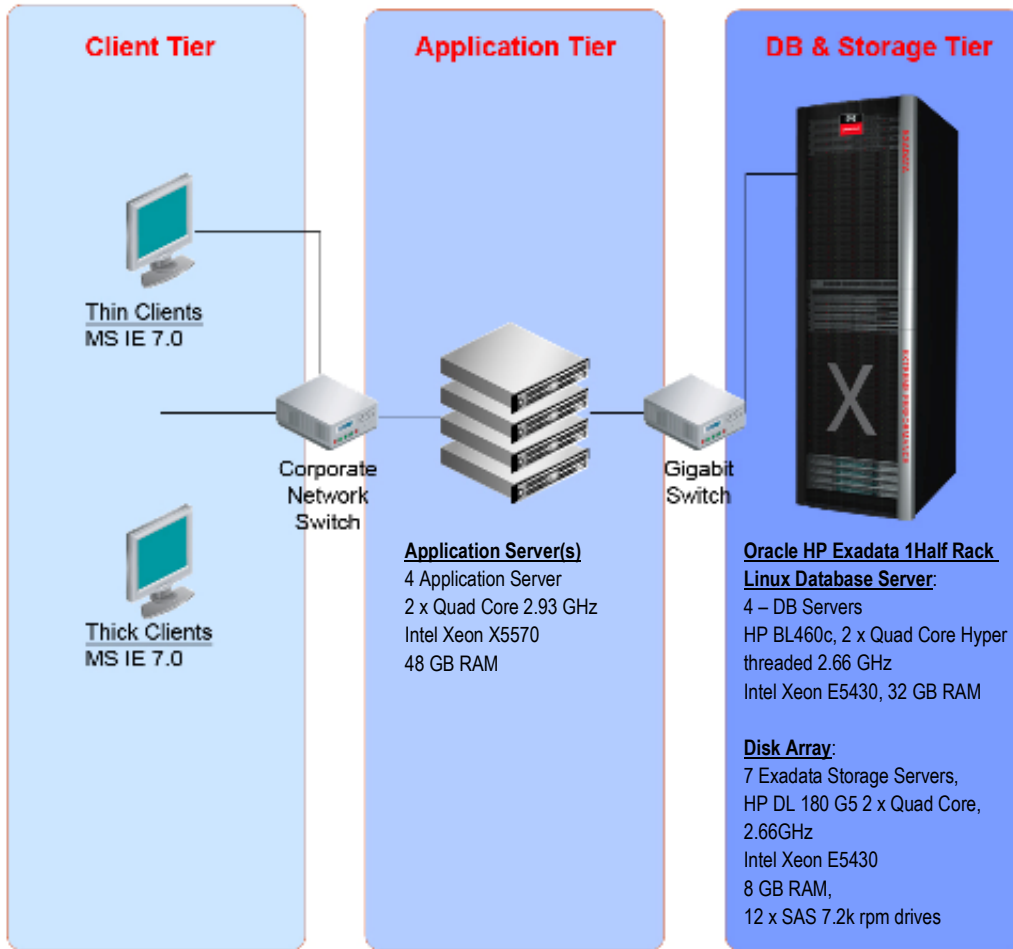
VALIDATION	NUMBER OF ROWS	NUMBER OF ROWS PERCENT
Interval Cuts to perform VEE	5,000,000	100%
Interval Validation Spike Check	100,000	2%
Interval Validation Gap Check and Estimation	50,000	1%
Interval High Low Failure	200,000	4%
Replacement Usage	100,000	2%
Interval Validation Meter Multiplier	100,000	2%

## Workload Profile

The business scenario tested in this benchmark was the key batch processing in Meter Data Management, i.e. Interval read storage, validation and Load.

## Performance and Benchmarking Environments

The benchmark was executed on Exadata 1 Hardware. The environment consisted of four servers in each tier (i.e. application & database tiers). On the database tier four servers represent four RAC nodes. On the application tier four servers contain the Oracle Utilities Meter Data Management version 1.5 application, which uses the Adapter Services in order to load the test data. Exadata storage server comes with an option of SATA or SAS disks. The benchmark environment had SATA disks. SAS disks have 1.5 times IOPS capacity relative to SATA. Thus adding 30 SATA spindles (total of 114 SATA spindles) is equivalent to 84 SAS disks. After running two application servers, the 30 additional drive resources were added to maximize the IOPS capacity available on the disks. The additional drives were used when using three and four application servers in the remainder of the horizontal scaling tests, below is a high-level architecture topology diagram.



Overall Hardware Topology for Exadata Benchmark

### Exadata Configuration

For this benchmark, the main table for storing usage was partitioned, and the LOB data for that table was stored using Oracle Secure Files with the compression level set to medium. The following table further defines the specifications of the database and the host servers.

TABLE 2. DATABASE SERVERS/HOSTS	
DATABASE ATTRIBUTE	ATTRIBUTE VALUE
Platform	HP BL460c
Operating System	Linux
O/S Version and Release	2.6.18-128.el
Cluster Software / Version	Oracle Clusterware 11.1.0.7

Number of CPUs	2 x Quad Core
Processor /CPU Speed	Intel Xeon E5450 / 2.66GHz
Memory	32GB
Storage System	7 Exadata Storage Servers + additional 30 SATA spindles
Raid Level	RAID10

## Application Servers

The following table outlines the application operating system and hardware specifications. One application server was used for the vertical scalability testing, two application servers were used for the horizontal scalability testing and data loading, and three application servers were used for the bill determinant calculation.

**TABLE 4. APPLICATION/BATCH MIDDLE TIER**

Platform	Intel Xeon X5570
Operating System	Windows 2003 R2 64-bit Server
O/S Version and Release	N/A
Software / Version	Oracle Utilities Meter Data Management version 1.5.0.6.0
# CPU	2 socket 4 core
Processor / CPU / Speed	2.93GHz
Memory	48Gb

## Result

With the hardware and application configurations outlined in this document, five million interval reads (i.e., one day of 15 minute interval data for 5 million meters) were validated and loaded in 25 minutes.<sup>1</sup> At the highest level of scaling a new maximum measured throughput was reached for the key use case (Interval Read Load). The maximum throughput reached 3328.9 transactions per second.<sup>2</sup>

Scalability within a single application server (vertical scalability)

Vertical scalability of Oracle Utilities Meter Data Management version 1.5 was demonstrated by a near-linear increase in application throughput with the increase in hardware resource utilization (in most cases CPU utilization).

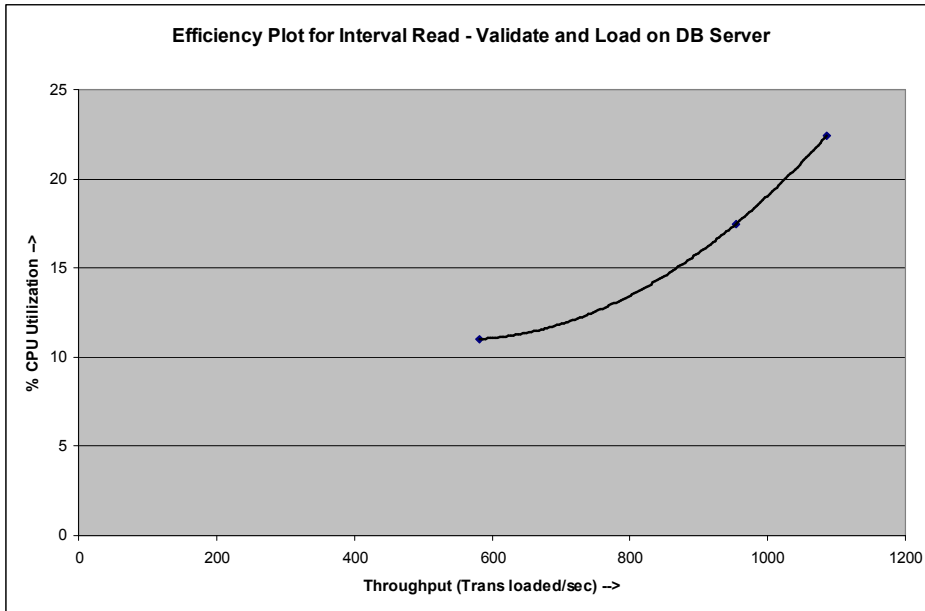
Configuration for vertical scalability:

- The first plot consists of 1 application server running 16 Oracle Utilities Adapter runtime services.
- The second plot consists of 1 application server running 32 Oracle Utilities Adapter runtime services.
- The third plot consists of 1 application server running 48 Oracle Utilities Adapter runtime services.

---

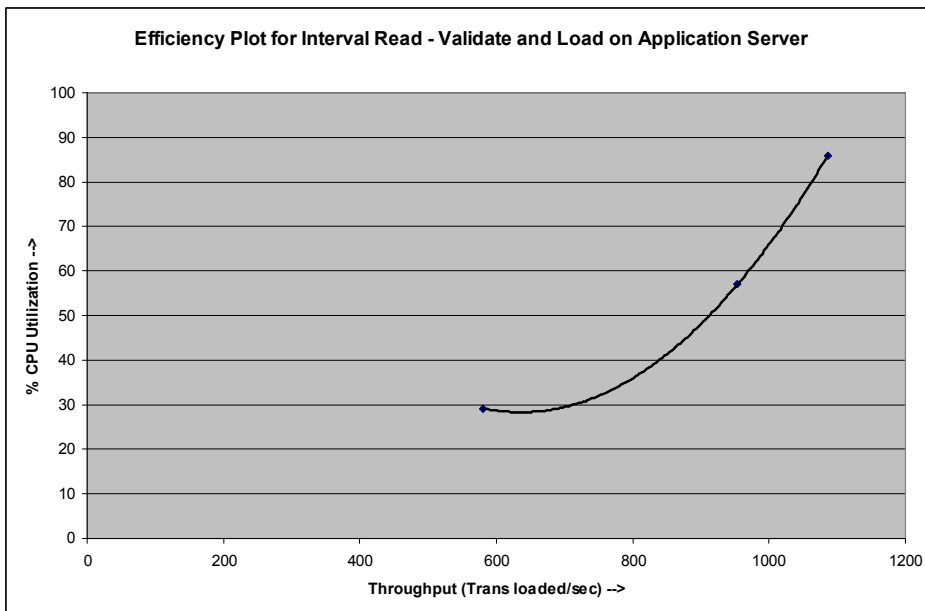
<sup>1</sup> Actual results may vary, based on a broad range of implementation-specific factors, such as transaction mix, hardware platform, network parameters, and database size. Oracle does not warrant or guarantee that customers will obtain the same or similar results, even if they use the same or similar equipment and/or software applications. Oracle does not warrant, endorse, or guarantee any performance of any products, any results desired or achieved, or any statements made within this document.

<sup>2</sup> Actual results may vary, based on a broad range of implementation-specific factors, such as transaction mix, hardware platform, network parameters, and database size. Oracle does not warrant or guarantee that customers will obtain the same or similar results, even if they use the same or similar equipment and/or software applications. Oracle does not warrant, endorse, or guarantee any performance of any products, any results desired or achieved, or any statements made within this document.



Scalability chart for interval read validate & load for database server CPU%.

Throughput shown here is "Interval Read Validate & Load" per sec, where 1 interval read corresponds to 1 days worth of interval data for 1 interval meter collected at 15 minute frequency. CPU utilization shown here is average CPU utilization per node, on a 4 node half rack Exadata 1 cluster.




Scalability chart for interval read validate & load for application server CPU%.

CPU utilization shown here is average CPU utilization on a single application server.

Scalability through multiple application servers (horizontal scalability)

Configuration for horizontal scalability:


- The first plot consists of 1 application server running 48 Oracle Utilities Adapter runtime services.
- The second plot consists of 2 application servers running 96 Oracle Utilities Adapter runtime services.
- The third plot consists of 3 application servers running 144 Oracle Utilities Adapter runtime services.
- The fourth plot consists of 4 application servers running 192 Oracle Utilities Adapter runtime services.



Horizontal scalability chart for interval read validate & load for database server CPU%.

Throughput shown here is "Interval Read Validate & Load" per sec, where 1 interval read corresponds to 1 days worth of interval data for 1 interval meter collected at 15 minute frequency. CPU utilization shown here is average CPU utilization per node, on a 4 node half rack Exadata 1 cluster.

Maximum throughput was reached running 64 services per application server (not represented in the Figures below). This totaled 256 services. The max throughput was 3328.9 transactions per second. The 5 million interval cuts were loaded in 25 minutes.<sup>3</sup>



Horizontal scalability chart for interval read validate & load for application server CPU%.

CPU utilization shown here is the sum of the average CPU utilization for each server.

---

<sup>3</sup> Actual results may vary, based on a broad range of implementation-specific factors, such as transaction mix, hardware platform, network parameters, and database size. Oracle does not warrant or guarantee that customers will obtain the same or similar results, even if they use the same or similar equipment and/or software applications. Oracle does not warrant, endorse, or guarantee any performance of any products, any results desired or achieved, or any statements made within this document.

## Conclusion

This benchmark demonstrated that the Oracle Utilities Meter Data Management version 1.5 was able to validate & load 1 day's worth of interval reads (15 min frequency) for 5 million interval meters within 25 minutes, on a half-rack Exadata 1 environment.<sup>4</sup> Oracle Utilities Meter Data Management version 1.5 application's linear scalability was proven both vertically and horizontally. Linear horizontal scalability was observed in the case of going from one application server to two application servers. Horizontal scalability was tested until the point of being CPU bound on the db server was reached. This occurred in the test of running four application servers. Following is the summary of application scalability characteristics, observed during this benchmark.

- Oracle Utilities Meter Data Management version 1.5 demonstrated near linear Vertical & Horizontal scalability, until reaching max hardware capacity (Disk I/O per second) of half-rack Exadata 1.

The Oracle Utilities Meter Data Management version 1.5 demonstrated near linear vertical & horizontal scalability up to max CPU capacity of half-rack Exadata 1 database server nodes, with additional 30 SATA spindles (in addition to 84 spindles in a half-rack). With this proven linear scalability, Oracle Utilities Meter Data Management on Exadata easily supports the growth of data that results from smart meter deployments. This scalability is simply achieved by adding more Exadata rack and application servers. In addition, Exadata provides fault tolerance features through built-in redundancies in storage layers without affecting application performance.

---

<sup>4</sup> Actual results may vary, based on a broad range of implementation-specific factors, such as transaction mix, hardware platform, network parameters, and database size. Oracle does not warrant or guarantee that customers will obtain the same or similar results, even if they use the same or similar equipment and/or software applications. Oracle does not warrant, endorse, or guarantee any performance of any products, any results desired or achieved, or any statements made within this document.



Oracle Utilities Meter Data Management 1.5  
Performance Benchmark on Exadata 1  
July 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



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, 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. 0710

**SOFTWARE. HARDWARE. COMPLETE.**