

Benchmark: Oracle Healthcare Transaction Base on Intel® Xeon® Processors on the IBM System x3950

*An Oracle White Paper
August 2008*



Benchmark: Oracle Healthcare Transaction Base on Intel[®] Xeon[®] Processors on the IBM System x3950

ABSTRACT

Performance benchmarking of the Oracle Healthcare Transaction Base (HTB) 5.3 occurred at the Intel Solution Services Center in Chantilly, Virginia, in early 2008. These tests documented Health Level 7 (HL7) Version 3¹ (V3) messaging throughput and application programming interface (API) transaction response times associated with increasing the number of users and stored data volumes, and changing hardware configurations. Findings illustrate that improved response times can be achieved on cost-effective hardware platforms while applying substantial system loads.

INTRODUCTION

The increasing volume of healthcare information – especially with the recent interest in large, regional healthcare information exchanges – has made cost-effective, high-performance healthcare data platforms an imperative for the industry. The Intel[®] Digital Health Group, along with Oracle, IBM, and NetApp, tested a scalable architecture that allows organizations to support a healthcare data exchange deployment. Hardware infrastructures from Intel, IBM, and NetApp along with healthcare software components from Oracle can scale from small community organizations to large regional delivery networks, as well as support incremental additions in scale over time. These tests performed in 2008 extend the testing on smaller hardware configurations on Oracle's Healthcare Transaction Base (HTB) 5.2 performed in 2006.²

This 2008 performance benchmarking documents the performance characteristics of HTB 5.3 based on increasing the number of users and message size. The following tables present the configurations and results of these benchmarking tests.

KEY BENCHMARK OBJECTIVES

- Demonstrate HTB's performance characteristics for large-sized healthcare IT systems with many patients and users
- Determine HTB's scalability characteristics with increasing the number of users (increasing transaction volume)
- Determine response time characteristics with large data volumes
- Determine maximum message throughputs
- Determine HTB's performance characteristics with increasing message size

¹ HL7 Version 3 refers to HL7 standards that are based on the HL7 Version 3 Reference Information Model. See www.hl7.org.

² "Benchmark: Healthcare Transaction Base on Intel[®] Xeon[®] Processors", an Oracle white paper, September 2006.

JOINT TESTING PROGRAM

- Oracle Software
- IBM Servers
- NetApp Data Storage

LOCATION AND DATE

- Intel Solution Services
- Chantilly, Virginia, USA
- April 2008

TEST VARIABLES

- Read/Write API Transactions (users)
- Messaging Throughput
- Message Size

Benchmarking Configurations

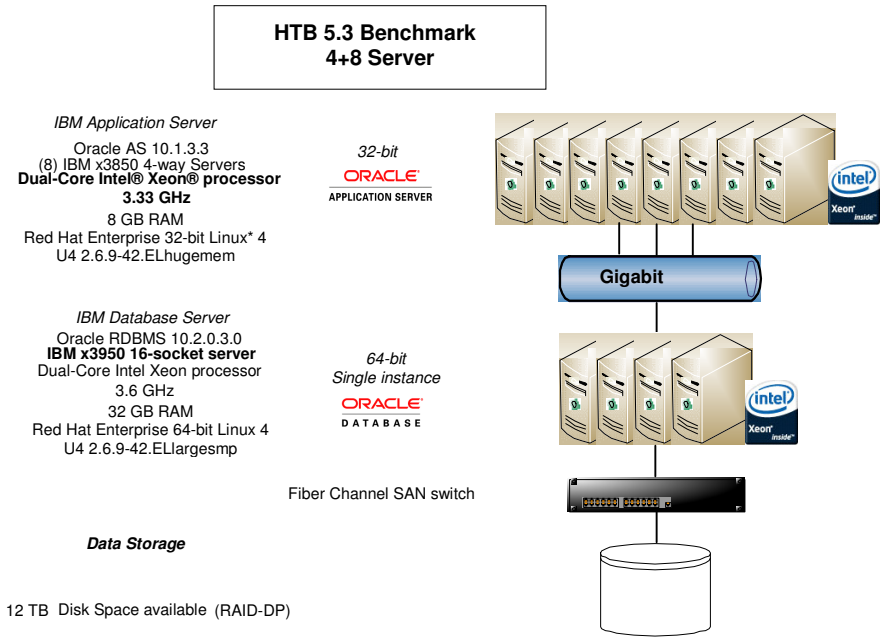


Table 1. Software Versions

<u>Component</u>	<u>Version</u>
Oracle Healthcare Transaction Base	V5.3
Oracle RDBMS (64-bit single instance)	10.2.0.3.0
Oracle Application Servers (32-bit)	10.1.3.3
Oracle E-Business Suite	11.5.10.2
Mercury LoadRunner*	8.1

Table 2. Software Configuration Detail (4 + 8 Server Configuration)

Oracle Database Servers (Total of 16 CPUs)	IBM x3950 with 16 sockets using NUMA architecture with (4) 4-socket chassis (supporting one 64-bit database instance) with Dual-Core Intel® Xeon® processor 4-way 3.6 GHz, 32 GB RAM Red Hat Enterprise 64-bit Linux* 4 U4 2.6.9-42.ELlargesmp
Oracle Application Servers (Total of 32 CPUs)	Eight IBM x3850 servers (supporting 32-bit software) with Dual-Core Intel Xeon processor 4-way 3.33 GHz, 8 GB RAM Red Hat Enterprise 32-bit Linux 4 U4 2.6.9-42.ELhugemem
Storage	12 TB SAN available (about 65 percent full with about 8 TB of data files)
Unique Patients	8 million
Data per Patient	1.0 MB
Years of Data	Dates ranged over 3 years (one inpatient encounter plus five outpatient encounters in 3 years)
Encounters	48 million (inpatient encounter = 360 clinical acts; outpatient encounter = 24 clinical acts)
Total Act Instance Count	3,888,000,000 acts
Average Message Load	50,000 HL7 messages/hour (average message size 25 KB with messages ranging up to 1 MB)

Table 3. IBM x3950 Database Server

Processor Sockets	16 processor sockets using NUMA architecture with (4) 4-socket chassis (supporting one 64-bit database instance) with Dual-Core Intel® Xeon® processor 4-way 3.6 GHz Red Hat Enterprise 64-bit Linux 4* U4 2.6.9-42.ELlargesmp
Memory	32 GB per chassis in this proof point, upgradable to a maximum of 512 GB
Chipset	XA-64e* Third-Generation Chipset
Cache in Chipset	Xcel4v* Dynamic Server Cache
Internal Storage	Six 2.5-inch SAS hot-swap hard drive bays
Size	3U for each building block

Table 4. IBM x3850 Database Server

Processor Sockets	4 processor sockets (supporting 32-bit software) with Dual-Core Intel® Xeon® processor 4-way 3.33 GHz Red Hat Enterprise 32-bit Linux 4* U4 2.6.9-42.ELhugemem
Memory	8 GB in this proof point, upgradable to a maximum of 64 GB
Chipset	XA-64e* Third-Generation Chipset
Cache in Chipset	Xcel4v* Dynamic Server Cache
Internal Storage	Six 2.5-inch SAS hot-swap hard drive bays
Size	3U for each building block

Table 5. Data Storage Detail

NetApp FAS3070	Enterprise storage solution for its mix of performance, scalability, and lower total cost of ownership. Scales to more than 500 TB of raw storage.
Product Details	RAID-DP* (high performance RAID 6) FC—28 (26 data disks plus 2 parity disks) SATA—16 (14 data disks plus 2 parity disks) FC—14 (13 data disks plus 1 parity disk) Qty (144) 144 GB FC disk drives; 15,000 RPM
Configuration Used	12 TB Available Disk, RAID-DP
SAN Protocol Support	Fiber Channel Protocol (FCP) for SCSI (2 Gb and 4 Gb); fabric-attached and direct-attached; iSCSI
Supported Configuration	Highly Available Controller Configurations: Active/active controller with active/active controller, active/active controller with stretch (non-switch) MetroCluster, active/active controller with fabric-attached MetroCluster Highly Available Back-End (Disk) Configurations: Dual-path, multipath HA storage
Maximum LUNs	Up to 2,048 front-end LUNs on FAS3040 and FAS3070
Management	Full-duplex 10/100/1000Base-T Ethernet onboard console, diagnostic LED/LCD, Remote LAN Management (RLM) Card, SNMP, telnet, SSH, HTTP, Web (SSL), host scripting, e-mail alerts

BENCHMARKING METHODS

USER LOAD

- **Users were added (via load generators) until transaction response times began to degrade. Results were reported at 75 percent of the peak number of users.**

User-Load Scalability: The number of users was increased on the system and key metrics were captured until the overall system throughput started to drop; that is, the system started running out of resources to support additional load. These loads were applied by increasing the number of application users (via load generators) simulated with Mercury LoadRunner*. Unless otherwise specified, results were reported at 50 to 75 percent of this high point for read/write API transaction throughput. This 75-percent benchmark was chosen because most production systems will try to keep peak concurrent users within 50 to 75 percent of capacity to keep performance high under all loads.

User Load Calculations: Different kinds of users in clinical care will generate user transactions at different rates. For example, transactions generated at the patient bedside while asking questions will generate fewer transactions than the review of a high volume of lab results at the nursing station. For this study, average clinical users were estimated to cluster about four transactions into one workflow at about one-minute intervals. Further, each average user would perform about thirty workflow sequences per hour, a reasonable assumption given that most clinicians in patient care need to interact with patients as well as the computer during a workflow sequence. Therefore, transaction volumes were equated with concurrent clinical users according to this formula:

$$\text{Concurrent Clinical Users} = (\text{x transactions / hour}) \div ((30 \text{ workflow sequences / clinical user / hour}) \times (4 \text{ transactions / workflow sequence}))$$

APPLICATION MIXES

- **Application mixes consisted of both read and write transactions.**

HTB supports a large library of read/write Java* APIs for application development. A variety of API transaction mixes simulated various types of application load (numbers in parentheses refer to the approximate ratio of user activity in the system at any given instant):

- **Application Mix 3 (API TMX):** Results retrieval (40 percent), clinical data entry (10 percent), order management (30 percent), patient registration (10 percent), EMPI (10 percent)

Each application mix was measured with a constant messaging load of 50,000 messages per hour on the system. These HL7 V3 messages averaged 25 KB in size and consisted of ADT, orders, labs, and public health messages.

Peak System Performance: To test the effect of processing different mixes of RIM classes, attributes, and data types, and to calculate the effect of message size on messaging throughput, additional tests were performed on a “messaging only” system load using these message type mixes and sizes:

- Message Mix 4 (MSG TMX): ADT 35 percent, lab results 25 percent, orders 25 percent, public health 15 percent
- Message Mix 4 Prime (MSG TMX Prime): ADT 30 percent, lab results 25 percent, orders 20 percent, public health 15 percent, PSM/CDA (100K) 10 percent
- Message Mix 5 (LRG MSG TMX): PSM/CDA 100 percent (large messages)
 - Three variations:
 - 100,000 messages
 - 500,000 messages
 - 1,000,000 messages

BENCHMARK RESULTS

User Load Scalability

The user load scalability tests demonstrate the effect of increasing transaction load on the response time.

USER TRANSACTION RATES

- Increasing rate of user read and write requests increases transaction completion rates up to a point.
- Further increases in user requests actually degrade transaction completion rates.
- Reporting 75 percent of peak user transaction completion rates allows a margin of safety in system configuration.

USER TRANSACTION RESPONSE TIME

- Measures return of data to user requests.
- Does not include graphical presentation time.
- Increases gradually with increasing users up to a point.
- Further increases in user requests rapidly degrade response time.
- Although average response time was 0.27 seconds, some large volume transaction types, such as "all lab results retrieval" for a single patient averaged 0.6 seconds.

75 Percent Load	API TMX	
Throughput (transactions/hour)	946,800	
Concurrent Users	7,890	
Response Time (seconds) *Key Results	0.27*	

$7,890 \text{ Concurrent Clinical Users} = (946,800 \text{ transactions / hour}) \div ((30 \text{ workflow sequences / clinical user / hour}) \times (4 \text{ transactions / workflow sequence}))$

Peak System Performance – Maximum Messaging Throughput

These tests demonstrate peak messaging throughput on a system without any other transaction load. Message size predicts message throughput.

MESSAGING THROUGHPUT

- Includes storing messages in logs.
- Includes transforming message data to relational format.
- Doubling message size decreases message throughput by about half in a linear relationship over a 40-fold increase in message size.

MSG TMX Prime	250,000 messages / hour	
LRG MSG TMX— 100,000	125,000 messages / hour	
LRG MSG TMX— 500,000	30,000 messages / hour	
LRG MSG TMX— 1,000,000	13,000 messages / hour	

DISCUSSION

HTB DESIGN CONSIDERATIONS

- HTB transforms messages to relational data

In the design of the HTB system, HL7 V3 messages are parsed and the data within the messages are written to a relational format that is optimized for transactional healthcare applications. Consequently, the native XML of the V3 messages is saved only in messaging logs destined for archive. Therefore these messaging throughput results include the CPU cycles required for data transformation and relational storage of discrete data in addition to the simple processing of message text into database storage. This detail of the design may be important when comparing these results to system designs that simply store the XML of the V3 messages.

DATA VOLUME MEASUREMENT

- Act Instance Count

This study uses the concept of HL7 V3 Act Instance Count as a measure of HL7 V3 data volume in performance testing. In these results the number of HL7 Encounter Acts and other HL7 Clinical Acts were reported. These two categories of HL7 Acts may be combined into a single measure of HL7 V3 data volume named “Act Instance Count.” This Act Instance Count is easily measured in messages and in databases of different designs as long as the instance identifiers are preserved. Act Instance Count, in addition to counts of unique patients and other V3 Entities, may allow better comparisons of benchmark performance metrics when comparing V3-based systems in the future.

INTERPRETING RESPONSE TIMES

- Average API response times
- End-user response times
- Response time ranges

API response times should be interpreted as representing only a portion of the response time the end user would experience in a graphical user interface (GUI) environment. If the GUI processing time added 0.5 seconds to an API response time of 0.5 seconds, the total response time the end user would experience would be 1.0 second. That said, an average API response time of 0.5 seconds can be considered “good” in most GUI environments. The API response times in these tests varied from <0.2 seconds to 1.0 second, depending on the volume of data returned.

TESTS AND PUBLISHED RESULTS

- Need for published results at large population sizes

Different kinds of healthcare activities require different numbers of transactions to complete specific tasks. In other words, typical workflow steps within healthcare require a sequence of transactions. To complete the steps in HTB, these workflow activities require these numbers of different transaction types:

- Results retrieval (6 total transactions)
- Order management (16 total transactions)
- Clinical documentation (5 total transactions)
- Patient registration/ADT (9 total transactions)
- EMPI (10 total transactions)

For example, these transactions are frequently used to complete results retrieval:

1. RES_PAT_LIST: Query for a patient list
2. RES_PAT_QRY: Query for selected patient
3. RES_LAB_QRY: Query for all lab results for a patient
4. RES_RAD_QRY: Query for all radiology reports
5. RES_ONE_RAD: Query for single radiology report detail
6. ORD_QRY_ALL: Query for existing orders

These testing center results are best interpreted in light of messaging volumes and user loads reported from actual medical centers. However, a review of the literature reveals no reports of messaging volumes and user loads in medical record collections approaching eight million active patients, as would be found in a large regional population.

It also should be noted that actual production environments probably experience wider peaks and valleys in load over days, nights, and weekends than those simulated in the laboratory environment. Better reporting in the literature would enhance the opportunity for realistic comparisons of performance testing.

**COMPARABILITY OF HL7 V2 AND V3
MESSAGE SIZE**

- No method for comparing data volume density in messages.

Another factor in evaluating HL7 V3-based laboratory testing includes comparing HL7 V3 laboratory results to today’s predominantly HL7 V2 production environments. The density of actual data in HL7 V2 “pipe and bar” messages is higher than that in HL7 V3 XML messages. Therefore actual data volume cannot be compared based on message size alone when comparing HL7 V2 and V3 messages.

Finally, benchmarking HTB performance on a larger hardware platform with more patients and more data allows some limited comparison to prior tests conducted on smaller hardware platforms in 2006. Although these are somewhat apples-to-oranges comparisons in that they are limited by the use of more modern and faster hardware, and more modern and faster software platforms, these comparisons do validate that the addition of application-level servers complemented by increased database server resources can dramatically improve the size of population supported by similar software.

Comparing 2008 with 2006

COMPARING 2008 WITH 2006

- Performance scales linearly with additional servers.

	<u>2006</u>	<u>2008</u>
HTB Version	5.2	5.3
Number of CPU in DB tier	2	16
Number of CPU in middle tier	12	32
Stored data volume	1TB	8TB
API Transaction Mixes:		
Max transactions/second (TPS)	22	263
Background messaging load	10,000 per hour	50,000 per hour
TPS per DB CPU	11	16
Maximum active users	660	7,890
Average user response time	0.50 seconds	0.27 seconds
Messaging Mixes:		
MSG TMX (2006)	31,200 per hr	272,000 per hr
MSG TMX prime (w/100Kmsg)	N/A	250,000 per hr
LRG MSG TMXa (100k)	N/A	125,000 per hr
LRG MSG TMXb (500k)	N/A	30,000 per hr
LRG MSG TMXc (1M)	N/A	13,000 per hr

CONCLUSION

- HTB scales with increases in stored data volume. Stored data volume was multiplied by 8 in 2008 over 2006 with no noticeable increase in response time.
- HTB scales with applied transaction/user load. As the results demonstrate, systems should be sized based on the response times observed between 50 percent and 75 percent of peak throughput in order to maintain a safety margin. However, as users are added or user intensity increases, it can be expected that transaction throughput and messaging throughput scales fairly linearly with increased application server and database processing power up to the volumes tested.
- Increasing or decreasing message size does not impact message-processing efficiency. This indicates that system administrators should base system requirement calculations only on total data volume, independent of message size.

Product Descriptions

Oracle Healthcare Transaction Base: HTB is a platform for the integration, development, and operation of healthcare applications across the healthcare continuum. It consists of a comprehensive, healthcare-specific data repository and standards-based information model coupled with a set of integrated services for enterprise terminology, data normalization, security and auditing, and business process and workflow. The platform is designed specifically for the healthcare industry and supports meaningful data consolidation, rapid application development, and genuine interoperability across different systems. With Oracle HTB, organizations in the business of healthcare – from hospitals and insurers to government bodies and public health agencies – can more effectively integrate, manage, deliver, and display information through the entire process of providing services. For more information, e-mail dan.russler@oracle.com or visit: <http://www.oracle.com/industries/healthcare/htb.html> http://www.oracle.com/industries/health_sciences/index.html

Intel® Xeon®: Intel Corporation's Dual-Core server platforms are based on the innovative 64-bit Intel® Xeon® processors. Combined with Fully Buffered DIMM (FBDIMM) technology, the Dual-Core architecture improves performance three times over Intel's previous Pentium® processors while cutting power use by 35 percent. Shared L2 cache allows each of the processors to share both code and data segments greatly increasing overall performance. The FBDIMM architecture streamlines memory access for database buffer blocks providing a 3X improvement over previous memory technology. The Intel® I/O Acceleration Technology delivers up to twice the data movement while cutting processor overhead by as much as 40 percent, supporting faster data transfer from database servers to requesting middle-tier applications. For more information, e-mail Stuart.A.Mathews@intel.com or visit:

<http://www.intel.com/technology/computing/dual-core/>
<http://www.intel.com/healthcare/>

ORACLE HEALTHCARE TRANSACTION BASE

The following features are included in HTB:

- Clinical data repository
- Enterprise terminology normalization tools
- Extensive read/write Java® API library based on the HL7 Reference Information Model

INTEL® DUAL-CORE SERVER ARCHITECTURE

The following features are included in Dual-Core:

- 64-bit Intel® Xeon® microarchitecture
- 35 percent power reduction over previous Pentium® processors
- Increased instructions per clock cycle
- Advanced smart cache
- Extended memory 64 technology
- Intel® I/O Acceleration Technology

IBM SYSTEM X3950

The IBM x3950 (database server) includes the following features:

- Third generation of IBM's Enterprise X-Architecture*
- Scales from 4 to 16 processor sockets in 4-socket building blocks
- Supports Dual-Core Intel Xeon processor MP
- Up to 512 GB memory
- Modular XpandOnDemand*
- Flexible system partitioning

IBM System x3950: The challenge that Oracle database customers are facing today is building an infrastructure that is highly available, yet scalable enough to meet the demands of a dynamic business environment. The x3950 is the ideal answer for customers choosing the Intel® processor-based platform for their Oracle implementations on Microsoft Windows* or Linux* operating system environments with the ability to scale up to 32-way processing capability.

The x3950 server allows customers to start with a single four-socket building block composed of Intel Xeon processors, memory, I/O cards, and hard drives. Then as needs grow, up to three additional four-socket building blocks may be added for true investment protection – unique in the commodity processor market.

IBM working in conjunction with Intel invested over \$100 million USD into a unique chipset designed to maximize the performance of the Intel Xeon processor 7000 series. The third generation of the Enterprise X-Architecture*, embodied in the X3 chipset, is designed to increase the movement of data from memory to each processor socket and to decrease the amount of socket-to-socket communication when a cache miss occurs. In fourth quarter 2007, this chipset was further improved with the release of the fourth-generation of Enterprise X-Architecture, resulting in the eX4 technology chipset. This new chipset supports the Quad-Core Intel® Xeon® processor 7300 series, as well as the new six-core Intel® Xeon® Processor 7400 series, and helps to drive twice the memory capacity of the previous generation. Clearly, the x3950 M2 with the eX4 chipset will help customers achieve even higher levels of system performance.

IBM SYSTEM x3850

The following features are included in the IBM x3850 (application server):

- Third generation of IBM's Enterprise X-Architecture
- Scale from 1 to 4 processor sockets
- Supports Intel Xeon Dual-Core Processor MP
- Up to 64 GB Memory

IBM System x3850: This is the entry-level server for the X3-based family. The x3850 is a single chassis server with up to four Intel Xeon processors 7000 series and up to 64 GB of memory. As with the x3950 server, a newer model based upon the x3850 M2 has been released based upon the new eX4 technology.

For more information visit:

<http://www-03.ibm.com/systems/x/hardware/enterprise/index.html>

NETAPP FAS3070

The following features are included in the NetApp FAS3070:

- Provides shared access to UNIX*, Windows, Linux, and Web data
- Simultaneously supports Fiber Channel SAN, IP SAN (iSCSI), and NAS
- Serves data at higher than 99.99+ percent availability

NetApp FAS3070: NetApp offers a comprehensive suite of award-winning products, software, and services that helps to simplify complex healthcare environments and dramatically reduce total cost of ownership. NetApp's fabric-attached storage (FAS) and inexpensive SATA systems integrate easily into complex healthcare environments and provide shared access to UNIX*, Windows, Linux, and Web data while simultaneously supporting Fiber Channel SAN, IP SAN (iSCSI), and NAS. These high performance systems have a proven ability to continuously serve data at higher than 99.99+ percent availability.



Benchmark: Oracle Healthcare Transaction Base on Intel® Xeon® Processors
on the IBM System x3950
August 2008
Authors: Sanjoy Mondal and Dan Russler, MD

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 © 2008, 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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners. 0408