



Oracle Retail Advanced Inventory Planning v 12.0 on System p 570 with POWER6 processor technology

*Ralf Schmidt-Dannert, IBM ATS
Sebnem Jaji, IBM
James Lowman, Oracle
Brad Metzler, Oracle
Brian Smith, Oracle
ISV Business Strategy and Enablement
August 2007*



Table of contents

Abstract.....	2
Introduction	3
System p 570 with POWER6 technology overview	3
Outstanding performance	3
Virtualization breakthroughs	3
Easy, balanced growth.....	4
Designed for business-critical environments	4
Tested-hardware configuration	4
Oracle Retail AIP architecture.....	5
Workload characterization and test-case descriptions	6
Performance comparison and scalability	6
Performance improvement with AIX 64 KB pages	12
Tested-software configuration.....	13
Application tuning.....	13
Database layout.....	13
Summary	14
Appendix A: Resources.....	15
About the authors	16
Acknowledgments	16
Trademarks and special notices.....	17



Abstract

This white paper illustrates the application performance gains for Oracle Retail Advanced Inventory Planning (AIP) version 12.0.1 on IBM POWER6. We observed this performance improvement after migrating from IBM POWER5 processor-based System p5-570 servers to IBM POWER6 processor-based System p 570 servers. This paper targets technical presales audiences as well as enterprises that are interested in deploying Oracle Retail AIP application on IBM System p.

Introduction

In July, 2007, IBM® and Oracle conducted performance tests with Oracle Retail Advanced Inventory Planning (AIP) version 12.0.1 to determine the performance gains that clients can experience by migrating from IBM POWER5™ processor-based System p5™ p570 servers to the new IBM POWER6™ processor-based System p™ 570 servers. The same tests also demonstrated the scalability of the Oracle Retail AIP solution with respect to supply chains that involve large numbers of SKUs.

The Oracle Retail AIP application is a suite of modules that is designed to manage the supply chains for large retailers at the supplier, warehouse, store and e-commerce levels. The system couples time-phased replenishment and allocation algorithms to produce an actionable receipt plan over time. This plan is based on demand forecasts, replenishment parameters and inventory availability at the numerous supply points within the supply chain.

Oracle AIP v12.0.1 has numerous functional, as well as technical and performance, enhancements — compared to previous versions.

System p 570 with POWER6 technology overview

Oracle Retail AIP v12.0.1 was tested on a 16-core IBM System p 570 model. This IBM system offers outstanding performance, breakthroughs in virtualization technology and easy growth. The p570 is also designed for business-critical environments.

Outstanding performance

The System p 570 server offers POWER6 processor-based configurations that contain from 2 to 16 cores and perform at speeds of 3.5, 4.2 or 4.7 GHz in an easily expandable 4U (4 EIA units) rack-mount modular package. Each building block of the p570 accommodates two processor cards, each of which provides two 64-bit POWER6 processor cores, 8 MB of Level 2 (L2) cache and 32 MB of Level 3 (L3) cache. The p570 provides up to 192 GB of memory capacity per building block — or 768 GB per system.

The POWER6 processor technology includes many design features that contribute to the excellent performance of the p570 system. Simultaneous multithreading allows two application threads to run at the same time, helping to reduce the time needed to complete tasks. Hardware-based decimal floating-point capability helps improve performance for financial calculations.

Virtualization breakthroughs

The IBM System p virtualization technologies logical partitioning (LPAR), IBM Micro-Partitioning™ technology, virtual LAN (VLAN) and Virtual I/O Server (VIOS) software are key elements of the System p family — offering exciting possibilities for application consolidation and system utilization. The POWER6 technology, which serves as the foundation for the System p 570 model, extends those IBM virtualization opportunities with functions such as Live Partition Mobility and Integrated Virtual Ethernet. *Live Partition Mobility* is designed to move running partitions from one POWER6 processor-based server to another without application downtime, providing enterprises with more flexibility for balancing workloads and improving application availability during planned system outages. *Integrated Virtual Ethernet* capability is standard in every p570 and enables easy-to-manage sharing of the integrated, high-speed Ethernet-adaptor ports.

Easy, balanced growth

System p570 accommodates enterprise requirements both for rapidly changing resource requirements in on demand environments and for long-term business growth. To help enterprises respond quickly and easily to changing business needs, the p570 offers optional Capacity on Demand (CoD) capability, allowing for the installation of processor and memory resources that are initially inactive and that you can activate for use as needed. For long-term growth, the modular design of the p570 makes it possible to easily and cost-effectively add processor cores, memory and I/O features as business needs increase.

Designed for business-critical environments

The p570 has an extensive set of reliability, availability and serviceability (RAS) capabilities to complement its processing power, I/O expandability, virtualization technologies and growth-on-demand features. This makes the p570 an outstanding system for a wide range of requirements — including database, business processing, application serving and business intelligence (BI) workloads, as well as server consolidation. The p570 has a standard one-year, next-business-day warranty that is upgradeable to include 24x7 support.

Tested-hardware configuration

Table 1 shows a summary of the POWER6 processor-based System p 570 used in these tests.

System p 570 with POWER6 processors	
Model	9117-MMA
Quantity	One
Operating system	IBM AIX 5L™ 5.3 TL06
Processor complex	16-core 4.7 GHz POWER6 processor-based
Memory	125 GB available to LPAR
Hard disks (internal)	66 internal, 72 GB hard disks 15 KB RPM (only used for application code)
Hard disks (external)	IBM System Storage™ DS4500, ESS800
I/O subsystem	Eight fiber-channel adapters

Table 1. Summary of the IBM System p 570 with POWER6 processor-based technology

Table 2 shows a summary of the POWER5 processor-based System p5-570 used in these tests.

IBM System p5-570 with POWER5 processors	
Model	9117-570
Quantity	One
Operating system	AIX 5L 5.3 TL05 SP6
Processor complex	16-core 1.9 GHz
Memory	62 GB available to LPAR
Hard disks (internal)	12 internal 72 GB hard disks
Hard disks (external)	IBM Enterprise Storage Server® (ESS) 800
I/O subsystem	Four fiber-channel adapters

Table 2. Summary of the IBM System p5-570 with POWER5 processor-based technology

Oracle Retail AIP architecture

For the performance tests discussed in this white paper, ORACLE relational database management system (RDBMS) and Retail Predictive Application Server (RPAS) were deployed on a single system. Only batch tests were performed.

AIP is one of several integrated applications within the Oracle Retail Suite (see Figure 1). Viewed at a high level, the process across the Oracle Retail Suites is as follows:

1. Oracle Retail Demand Forecasting (RDF) provides a forecast of consumer demand. This data is available to Oracle AIP.
2. The AIP batch run produces a receipt plan that can be acted upon by using replenishment parameters that are maintained inside of Oracle AIP. A merchandising system, such as Oracle Retail Merchandising System (RMS), provides the hierarchy and inventory data.
3. The receipt plan then goes to the order-management system inside of Oracle AIP, where orders that await fulfillment are formally prepared for processing.
4. The order-management system then submits these orders to the merchandising system. These orders returned to Oracle AIP in a subsequent batch run as in-transit orders.
5. A collaborative planning, forecasting and replenishment planning (CPFR) application shares sales forecasts and orders with suppliers — so that trading partners can prepare for forthcoming orders.




Figure 1. Oracle Retail AIP software stack

Workload characterization and test-case descriptions

Replenishment and reconciliation were part of the AIP batch run that generated a receipt plan for stores and warehouses across the length of the planning horizon. The final output from this batch process was as follows:

- Constrained receipt plan for warehouses and stores across the fixed period
- Unconstrained receipt plan for warehouses and stores across the post fixed period through the end of the planning horizon

The workloads included in this benchmark consisted of the following test cases:

- Test case 1: 4000 SKUs – multidomain
- Test case 2: 8000 SKUs – multidomain
- Test case 3: 20 000 SKUs – multidomain

Eighty percent of the SKUs were planned for 35 days and twenty percent were planned for 175 days. See the following formulas:

$$(20\,000 \times 0.8) \text{ SKUs planned} \times 35 \text{ days} \times 1093 \text{ stores}$$

$$(20\,000 \times 0.2) \text{ SKUs planned} \times 175 \text{ days} \times 1093 \text{ stores}$$

The replenishment step is an AIP batch process that runs twice within the nightly batch. It calculates the replenishment plan for each SKU-pack size — at each store across the planning horizon.

The largest test case involved approximately 1.5 billion* planning opportunities.

Performance comparison and scalability

The “Tested-hardware configuration” section of this white paper shows the details of the hardware configurations that were used for the POWER5 and POWER6 performance comparison.

Figure 2 compares the POWER5 to the POWER6 processor-based system. The test cases differ in the number of items that were planned (4000, 8000 and 20 000 units, respectively). The largest test case involved the planning of approximately 1.5 billion decision variables — based on the number of SKUs and stores, as well as the number of days in the planning horizon.

Figure 2 shows good scalability up to 1.5 billion replenishment opportunities. The performance improvement varies for enterprises that migrate to AIP version 12.0.1 and to a POWER6 processor-based system. Depending on the results for the tested configurations, the expected performance improvement for the replenishment step is in the range of 32 to 47 percent.¹

***Note:** One billion refers to the America system billion, which is 1 000 000 000.




Figure 2. Elapsed time for the replenishment process on the POWER5 and POWER6 processor-based systems

The POWER5 processor-based system in test case 3 displayed throughput that was lower than expected when compared to the POWER6 processor-based system. During the execution of this test case on the POWER5 processor-based system, a significant I/O bottleneck was observed. This had the effect that the available processor resources were not fully used because the workload had to wait for the physical I/O to disk operation.

The key difference between the POWER5 and POWER6 processor-based systems was that the POWER6 environment was configured with twice the amount of physical memory (128 GB) and, therefore, was able to put a significantly higher percentage of the data that was accessed into cache. This resulted in significantly less physical I/O to disk operations. You can see this difference clearly in Figures 5 and 6.

Figures 3, 4, 5 and 6 show the performance behavior (processor utilization and I/O rates) of the POWER5 and POWER6 processor-based systems during the replenishment step for the workload used in test case 3.




Figure 3. Processor-utilization rate for replenishment on the POWER6 processor-based system




Figure 4. Processor-utilization rate for replenishment on the POWER5 processor-based system




Figure 5. I/O rates for replenishment on the POWER6 processor-based system




Figure 6. I/O rates for replenishment on the POWER5 processor-based system

Figures 7 and 8 show the scalability of Oracle Retail AIP on the POWER5 and POWER6 processor-based systems.




Figure 7. Multidomain test-case comparison for scalability on the POWER5 processor-based system




Figure 8. Multidomain test-case comparison for scalability on the POWER6 processor-based system

Performance improvement with AIX 64 KB pages

On the POWER6 processor-based System p p570, the effect of 64 KB page size on the Oracle Retail AIP application was measured. Figure 9 shows the performance gain that resulted by using 64 KB page data, stack and text pages for the RPAS processes. Note that the Y-axis does not start at zero.




Figure 9. Performance gain by using 64 KB page data, stack and text pages

Tested-software configuration

The combination of software that was tested for this white paper includes the following products:

- Application: Oracle Retail AIP v 12.0.1, RPAS 12.0.4
- Database: Oracle EE 9.2.0.7

Application tuning

Application-configuration modifications were implemented to enhance performance, including the following changes:

- The number of parallel execution threads was set to the number of local domains, so that all domains could be processed simultaneously.
- The commit values within the Oracle database were modified so that there were a small number of commits.
- The number of allowed-store source threads in the Oracle database was modified.

Database layout

Oracle Retail AIP is an I/O-intensive application that requires high I/O bandwidth between the system and the storage subsystem. Therefore, the data layout is a significant factor for better performance. The POWER6 processor-based p570 system was directly attached both to a DS4500 and an ESS 800 storage system, through eight fiber-channel adapters. The file systems for the RPAS data and the Oracle database were spread across the DS4500 and ESS 800 storage systems, as shown in Figure 10 and 11. The application and database files were placed in file systems in the volume-group that was named *sharedVG*. The DS4500 system had 91 disk units (seven drawers, each with 14 disk units), and the ESS 800 system had 144 disk units.




Figure 10. Distribution of application-data file systems on the DS4500 storage system




Figure 12. Distribution of application data file systems on the ESS800 storage system

The application and database executable codes were installed on the System p 570 internal disks.

Summary

When previous POWER5 processor-based p5-570 batch run times are compared to the POWER6 processor-based p570, a 32 to 46 percent reduction in the run time of the replenishment segment of the AIP batch process was observed. Some shorter-duration steps in the batch process show as much as a 57 percent reduction in execution time. This reduction in elapsed time enables enterprises to considerably shorten their window for the nightly-replenishment process.

Oracle AIP version 12.0.1 scales well up to 1.5 billion replenishment opportunities with respect to the size of the planning problem, as defined by the following calculation:

$$\text{number of SKUs} \times \text{number and locations} \times \text{length of planning horizon}$$

Appendix A: Resources

These Web sites provide useful references to supplement the information contained in this document:

- IBM System p Information Center
<http://publib.boulder.ibm.com/infocenter/pseries/index.jsp>
- IBM Publications Center
www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US
- IBM Redbooks®
www.redbooks.ibm.com/
- IBM System p 570
ibm.com/systems/p/hardware/midrange/570m/index.html
- IBM AIX PartnerWorld page
ibm.com/partnerworld/aix
- IBM POWER6
ibm.com/systems/p/about/power6.html
- ORACLE Retail applications
www.oracle.com/industries/retail/index.html
- ORACLE Retail Advanced Inventory Planning
http://www.oracle.com/application/retail/sc/adv_inv_plan.html

About the authors

Sebnem Jaji is a member of the IBM Systems and Technology Group. She works on benchmarking and optimizing independent software-vendor solutions for IBM Systems platforms. You can contact her at jaji@us.ibm.com.

Ralf Schmidt-Dannert is a member of the IBM Advanced Technical Support team and focuses on Oracle databases and the Oracle Retail solution for IBM System p servers. For questions and comments, you can reach him at dannert@us.ibm.com.

Brian Smith is a performance specialist from Oracle on contract within the Oracle Retail Performance Lab. Brian provides expertise on AIP and the underlying calculation engine.

Brad Metzler is a performance specialist within the Oracle Retail Performance Lab. He provides Oracle database support and expertise. Brad also analyzes and tunes database queries within Oracle Retail applications.

James Lowman is a performance specialist within the Oracle Retail Performance Lab. He provides benchmarking, performance analysis and AIP expertise.

Acknowledgments

The following Oracle Retail Performance team members contributed to this project:

Team member	Company	Role
Jessie Ge	Oracle	Performance specialist
Ron Goodson	Oracle	Project manager
James Lowman	Oracle	Performance specialist
Brad Metzler	Oracle	Performance specialist
Brian Smith	Oracle	Performance specialist

This project had the following sponsors:

Sponsor	Company	Role
Fred Unger	Oracle	Oracle Director of Retail Technology Group
Rick Wiederholt	Oracle	Oracle Director of AIP Software Development
Ramesh Chitor	IBM	Solutions Relationship Manager, Oracle Applications, Business Strategy and Enablement, Systems and Technology Group



Trademarks and special notices

© Copyright IBM Corporation 2007. All rights Reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

AIX, AIX 5L, Enterprise Storage Server, IBM, the IBM logo, Micro-Partitioning, POWER5, POWER6, Redbooks, System p5, System p and System Storage are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Benchmark results are highly dependent upon a number of factors, including workload, specific application requirements, and system design and implementation. Results under other conditions may vary significantly. IBM and Oracle make no warranty on system performance, price and price/performance and you should not rely upon this performance report including all information contained herein (this "Document") when making capacity planning or product purchase decisions.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.