



Tim Bohn  
Dr. Howard Operowsky  
Karen Lawrence

# Add Memory, Improve Performance, and Lower Costs with IBM MAX5 Technology

## Executive summary

The need for computer memory is increasing at a tremendous rate as databases continue to grow. With the increased use of virtualization, applications become memory-constrained and performance bottlenecks occur. At the same time, x86 processors are evolving with more cores and more thread capability. Without adequate memory, these processors are underutilized. This lack of memory is occurring more frequently and with more intensity as applications, databases, and technologies grow and change.

A previous solution was to purchase additional computers, and then distribute applications among them to use the combined memory capacity. However, the hardware, data center space, power and cooling requirements, and licensing were costly, and still the computers were underutilized.

IBM® is solving this issue with a new server technology in its IBM System x® eX5 class of server: the MAX5 memory expansion unit. This unit is an IBM-only, System x technology that adds memory apart from the processor. With no added processors, MAX5 is designed to increase productivity and reduce costs in a way that other servers of this class cannot.

At the IBM Strategy and Technology Lab (STL), we wanted to test how MAX5 improved system performance on web-based and database workloads. We used a production-quality IBM System x3690 X5 server, and we tested with and without MAX5. The results presented in this IBM Redpaper™ publication are from actual measurements taken during our STL testing.

In this paper, we demonstrate how the additional memory afforded by MAX5 improved the performance and costs of database and web-facing applications in our tests. In particular, our results show the following benefits:

- ▶ Response time for business intelligence reports was 1.5 - 2.8 times faster
- ▶ Cost of business intelligence reports was decreased by 31% - 64%
- ▶ Throughput of web-facing applications was 2.4 - 4.9 times greater
- ▶ Read/write response time was decreased by 60% - 80%
- ▶ It can cost 49% more for a web-facing application system without MAX5 to match the performance of a system with MAX5

This publication is directed to IT professionals and decision-makers, such as CEOs, CIOs, CFOs, information architects, business intelligence administrators, IBM DB2® database administrators, IBM clients, and IBM Business Partners.

## How MAX5 technology works

IBM System x MAX5 technology<sup>1</sup> is based on eX5, the next generation of Enterprise IBM X-Architecture®. The design of the MAX5 increases the number of dual inline memory modules (DIMMs) supported in the system up to 32. With MAX5 and 32 GB DIMMs, you can add as much as 1 TB of memory to an existing eX5 computer to support large, memory-constrained applications. This industry-first technology separates memory from the processor by using the following features:

- ▶ The advanced eX5 chipset:
  - The design of the MAX5 increases the number of DIMMs supported in the system up to 32. With MAX5 and 32 GB DIMMs, you can add as much as 1 TB of memory to an existing eX5 computer to support large, memory-constrained applications.
  - Attaches to QuickPath Interconnect (QPI) links as a node controller, thus giving the chipset direct access to all CPU bus transactions.
  - Maintains a directory of the last-level cache of each central processing unit (CPU), thus allowing the chipset to respond to memory requests before the end of a broadcast snoop cycle, improving performance.
- ▶ An additional 16 channels of memory bandwidth  
These added channels provide a significant boost to overall throughput and additional memory performance. Even on a system that is not memory-constrained, the additional channels allow the system to reach higher memory capacity by using smaller, lower-priced chips. The result is a more economical system.

As this paper shows, these benefits translate into faster response times, faster throughput, and lower costs.

---

<sup>1</sup> *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and BladeCenter HX5*, REDP-4650: <http://www.redbooks.ibm.com/abstracts/redp4650.html>

The System x3690 X5 with the MAX5 attached is shown in Figure 1.



Figure 1 IBM System x3690 X5 with MAX5 memory expansion unit attached beneath it

## About our products

Before we describe our tests, we look at some of the hardware and software components we used in the tests and describe why we chose them.

### IBM System x3690 X5 server

We used System x3690 X5 servers because of their inherent power. These servers are two-socket, rack-mount servers with up to 10 core Intel Xeon EX processors. Some x3690 X5 server models are combined with MAX5 (with its 32 GB of RAM), to supply up to 1 TB of memory.

This server has up to 16 hard disk drives (HDDs), supporting 6.35 cm (2.5 in.), hot-swap, internal Serial Attached SCSI/Serial Advanced Technology Attachment (SAS/SATA) drives, totaling up to 8 TB. The server also provides up to 24 solid-state drives (SSDs), supporting 4.57 cm (1.8 in.), hot-swap, solid-state drives, by using up to three IBM eXFlash storage options<sup>2</sup>. Each solid-state drive is designed to sustain approximately 4,300 input/output operations per second (IOPS). This IOPS rate is more than 10 times the rate achieved by an enterprise 15,000 RPM, 6.35 cm (2.5 in.), 146.8 GB SAS HDDs (420 IOPS).

### IBM MAX5 for System x

MAX5 for System x is a 1U memory expansion drawer. It adds 32 DIMM slots, the advanced eX5 chipset, and 16 additional channels of memory bandwidth, as explained in “How MAX5 technology works” on page 2. MAX5 is an optional unit for the IBM eX5 line of servers. It is also available as an option on all models of the System x3690 X5.

<sup>2</sup> IBM Unveils Industry's First Systems that Rewrite Economics of 'Industry-Standard' Computing, March 2, 2010: <http://www.ibm.com/press/us/en/pressrelease/29570.wss>

MAX5 for System x is scalable. In an IT environment, the base System x3690 X5 (without MAX5) with 16 DIMMs can serve as a starting point. When your IT needs expand, you can install an additional 16 DIMMs in your System x3690 X5. As your IT needs expand further, you can install the 1U MAX5 memory expansion unit, thus adding 32 DIMM sockets. The total number of sockets available to the system is then 64 DIMMs in only 3U of rack space. Using 16 GB DIMMs, your system has 1 TB of RAM.<sup>3</sup>

## DB2

DB2 is a family of IBM relational database management system (RDBMS) products known for a number of factors:

- ▶ The industry-leading database performance of DB2 is a significant reason for its use in analytical and online transaction processing (OLTP) workloads. Such performance is a factor in overall system performance. DB2 is also appropriate for mixed workloads on distributed systems.
- ▶ DB2 is scalable and as equally effective for single-user applications as it is in large, multiuser environments. Users can run structured query language (SQL) queries simultaneously on DB2.
- ▶ DB2 is efficient. With its application development and management tools, you can streamline database administrator tasks, detect emergent problems, offload inactive data, and improve the speed of backup and recovery processes. This efficiency translates to lower operational costs.
- ▶ DB2 supports multiple operating systems, including Microsoft Windows, Linux, UNIX, and IBM z/OS®. Administration tasks can be carried out at a command line or the GUI, which is a multi-platform Java client with multiple wizards for further increasing efficiency.

## IBM WebSphere eXtreme Scale

WebSphere eXtreme Scale is a low-latency, elastic, in-memory data grid that, to an application, looks like single, virtualized, in-memory cache. We chose IBM WebSphere® eXtreme Scale for our test because it takes full advantage of the additional memory provided by MAX5. A WebSphere eXtreme Scale grid is used effectively as a data cache for a database or other data source, which is generally slower to respond to requests due to the need to access data on a hard disk.

The elastic scalability of WebSphere eXtreme Scale is possible by using distributed object caching. The grid is implemented by combining the heap spaces of a collection of Java virtual machines (JVMs), which WebSphere eXtreme Scale calls *containers*. These containers can be located in the same physical machine, or they can be dispersed geographically among different physical machines. The data objects in the grid are partitioned among the heaps of the JVMs. Applications are not affected by which JVM contains the data because the grid is presented as a single entity. The grid can grow or shrink by adding or removing JVMs. Data in the grid can be replicated to enhance reliability.

In an OLTP setting, WebSphere eXtreme Scale can process billions of transactions a day. It does so by managing data and business logic across multiple servers and virtualized environments. WebSphere eXtreme Scale also delivers predictable response time.<sup>4</sup>

<sup>3</sup> *IBM System x3690 X5 at-a-glance guide*, January, 2012:  
<http://www.redbooks.ibm.com/abstracts/tips0818.html>

<sup>4</sup> *WebSphere eXtreme Scale Best Practices for Operation and Management*, SG24-7964:  
<http://www.redbooks.ibm.com/abstracts/sg247964.html?Open>

## About our tests

Two series of tests were conducted with and without MAX5: one tested database system performance; the other tested web-facing transaction performance. This section describes the overall purposes of the tests and their setups, configurations, procedures, and results.

### Tests for database system performance

In this test, our goal was to determine the impact of MAX5 on performance. This test followed a serial execution test in which successive reports were generated. The cost and performance of the system with and without MAX5 also were compared.

For these tests, we used a production-quality System x3690 X5 server and IBM DB2 to determine the effects of MAX5 on a 1.3 TB data warehouse in a business analytics environment.

A simplified view of a System x3690 X5 connected to a MAX5 is shown in Figure 2.

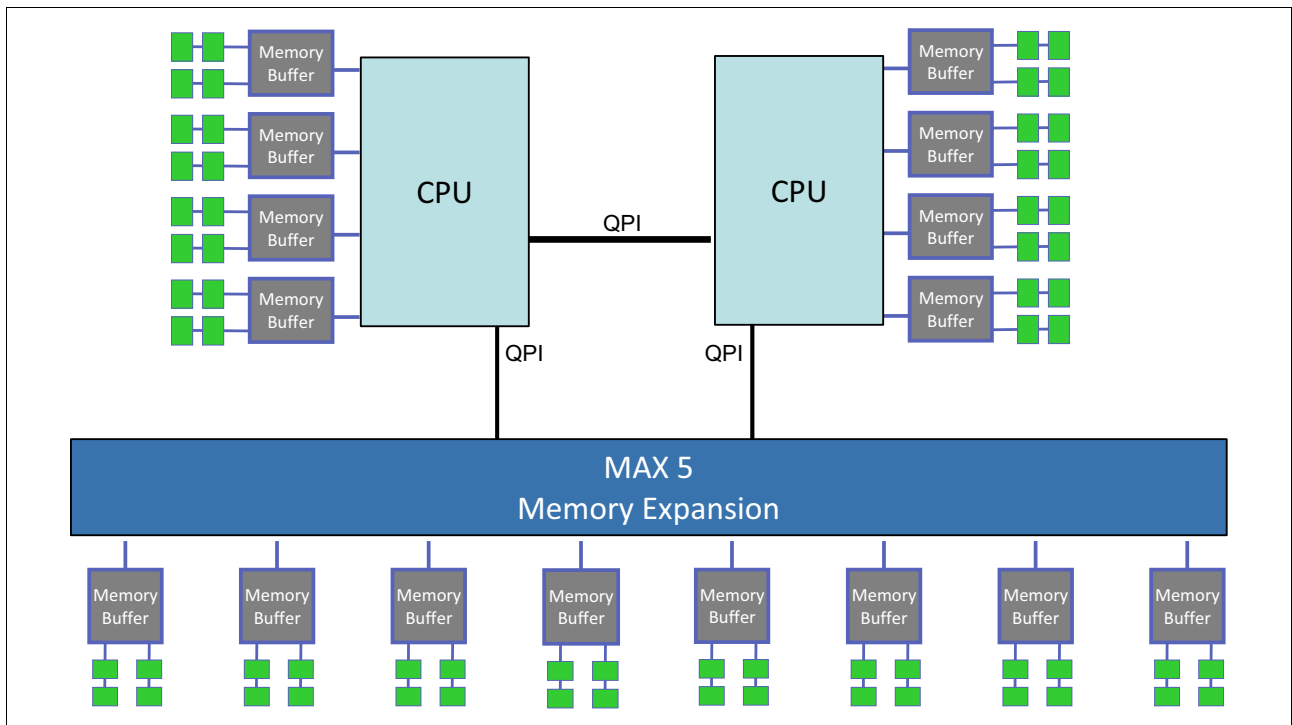


Figure 2 A System x3690 X5 connected to a MAX5

### Test setup and configuration

The server used in the test was an IBM System x3690 X5, with two, eight-core, Intel Xeon X7560 processors, for a total of 16 cores. Two memory configurations were used:

- ▶ With MAX5: 512 GB RAM internal to the System x3690 X5, plus 512 GB RAM installed in the MAX5 and attached to the System x3690 X5, for a total of 1 TB RAM
- ▶ Without MAX5: 512 GB RAM internal to the System x3690 X5

We used HDDs on the server for the operating system, applications, database logs, and database temp space. An IBM storage controller with 32 HDDs was attached for the database. This configuration provided plenty of storage I/O throughput. The database was a 1.3 TB data warehouse that used IBM DB2 V9.7.4 and ran Red Hat Enterprise Linux V5.6.

For the driver, a client machine played the role of a user that sent report requests in the form of sets of SQL queries. It was used to send the report requests and record the amount of time taken for the database server to complete the instruction. The hardware and software configuration used in our testing is shown in Figure 3.




Figure 3 Test hardware and software configuration

## Business Intelligence Day

We created an environment that simulated a business intelligence (BI) workload, called *BI Day*, which consisted of SQL queries from IBM Cognos® reports. These SQL queries are submitted by Cognos to the database management system (DBMS) to generate a report. Each report consisted of one to five SQL queries. The reports were designed to fall into these typical reporting scenarios:

- ▶ Simple: Minimal, fast-running reports, dashboards, and ad hoc reports. The reports take less than one second to run.
- ▶ Intermediate: Advanced reports that require predicate (a *where* clause) evaluation over a large fact table, then join an aggregation of a relatively small result set (only a small portion of the records are retrieved). These reports take from one second to one minute to run.
- ▶ Complex: Expert and resource-intensive reports, requiring multiple joins and aggregations on the full fact table (every row must be retrieved and operated on). These reports take 15 - 60 minutes to run.

We ran the first series of tests with the MAX5 attached and 1 TB total RAM. With this configuration, nearly all of the data warehouse (about 97%) was buffered in memory.

We ran the second series of tests without MAX5 and used the standard memory configuration of 512 GB. Here, more data than can be buffered in memory was present. The result was disk I/O whenever the data was not present in memory (about 20% of the time) and needed to be retrieved from disk. Both test series were run over the course of one week.

Shown in Figure 4, the serial execution test simulates a user who executes each of these reports consecutively and records the actual response time.

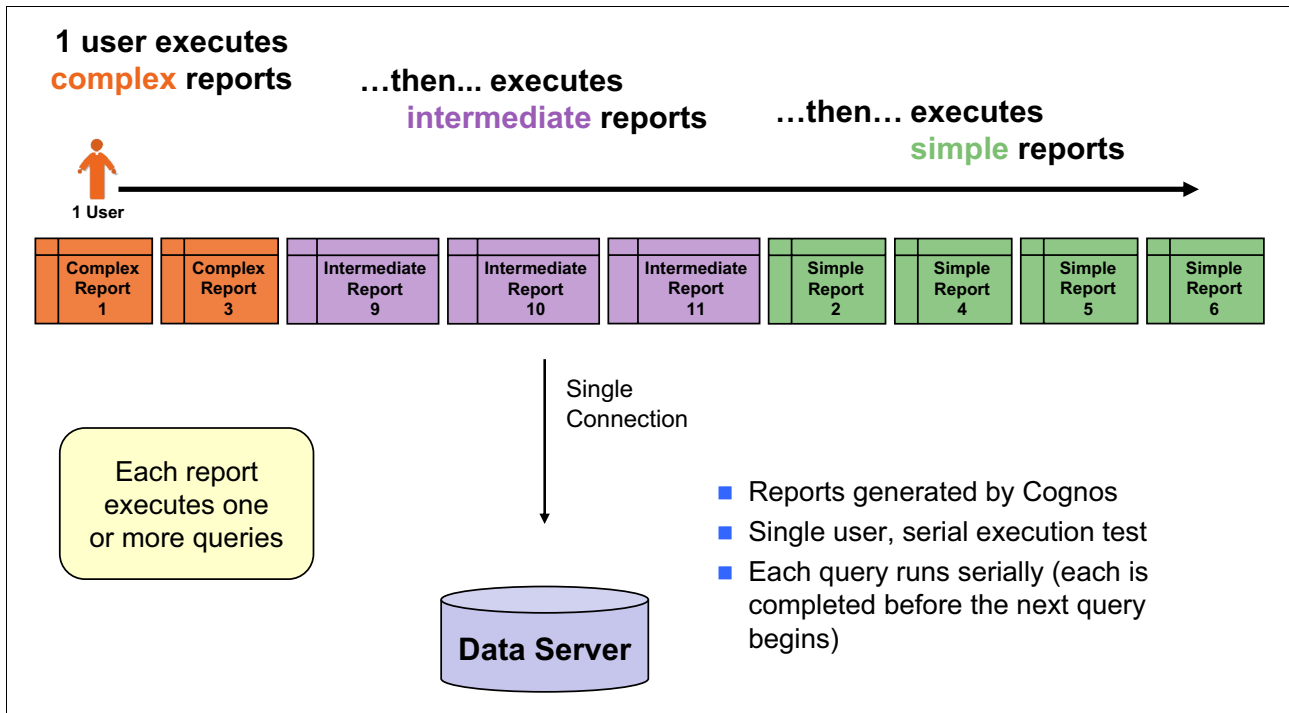


Figure 4 The serial execution test measured report performance for individual reports

## Results

The test results demonstrated that MAX5 can cache up to twice as much information in memory, thus reducing disk I/O and increasing performance and productivity. Having available extra memory for a business analytics and data warehouse workload was valuable. Response time was improved by caching the data in memory rather than continually reading from disk I/O.

With MAX5 for business intelligence workloads, you can produce twice as many reports, as we showed in our test over a one-week period (by using actual results). Therefore, over a three-year period, the cost of reports can be 31% - 64% lower with the addition of MAX5 (by using extrapolated results).

### **MAX5 turns reports around faster**

As shown in Figure 5, with MAX5 attached, and when the response times for the three report types are averaged, reports were completed 1.5 - 2.8 times faster. Of course, your performance numbers might vary based on workload profiles.

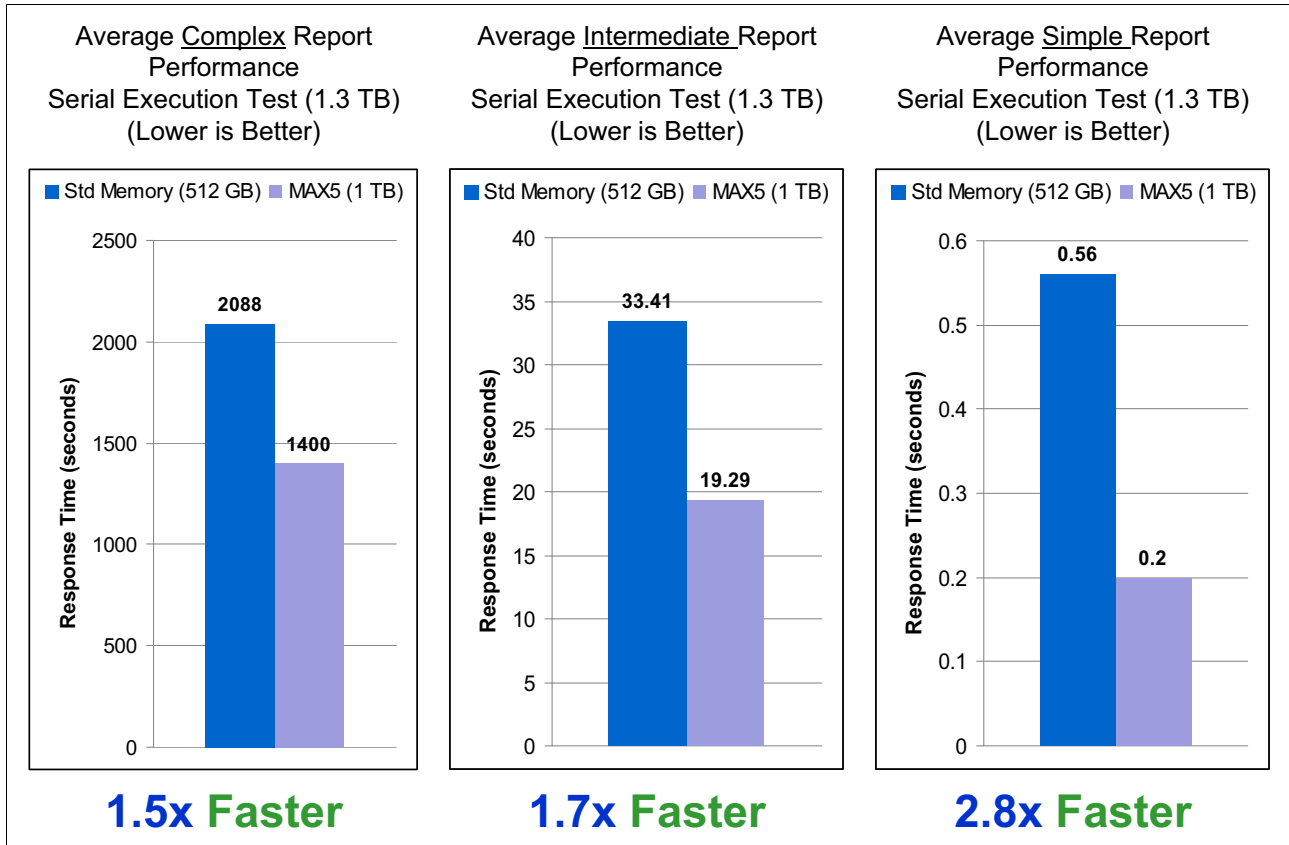


Figure 5 Comparison of report response times: MAX5 (1 TB) versus standard memory (512 GB)

What this means to you as the company that needs to produce reports is clear: in a one-week period, an organization can produce 432 complex reports with MAX5, but only 290 such reports without MAX5. Add intermediate reports and many simple reports and you can see that the company can generate twice as many reports in one week with MAX5. Without MAX5, your organization needs another complete server setup, including software licenses for additional cores, and data center space. Additional cooling and power needs also increase operating costs.

Suppose you have the hardware and software we used in this test and you need to expand. You can choose either option to meet your needs:

- ▶ Purchase an additional server and the software, software licensing, maintenance, data center space, cooling, and power for the server.
- ▶ Purchase MAX5 and the extra 512 GB RAM. With this option, you do not incur expenses for an additional server (hardware, software, software licensing, maintenance, data center space, cooling, and power). Rather, you incur expenses for the MAX5 drawer and 512 GB RAM only.

The difference in the results of the options is dramatic: the first option doubles your cost, a 100% increase; the second option increases your expenses by only 3%.



For more information about the cost savings achieved when using MAX5, see “MAX5 reduces the total cost of acquisition” on page 12.

**MAX5 lowers the cost of reports**

Although MAX5 decreases the time needed to generate reports, the question remains whether it is worth the investment. Consider the total cost of hardware and software over a three-year period, and the total number of reports that can be generated. The cost of reports is 31% - 64% lower for the MAX5-outfitted server. The difference in the cost of reports over that three-year period is shown in Figure 6. Your performance numbers might vary based on workload profiles.

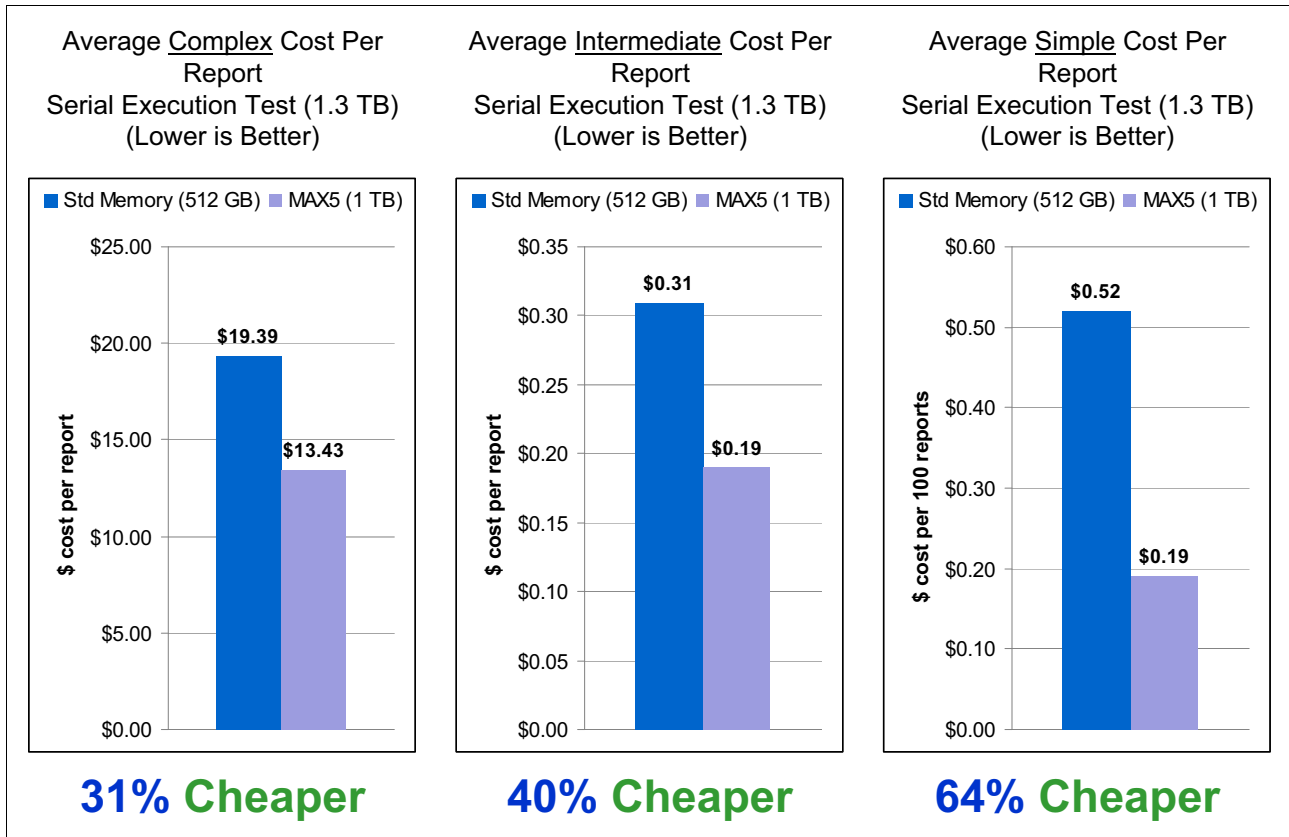


Figure 6 Comparison of cost of reports: MAX5 (1 TB) versus standard memory (512 GB)

You can see from Figure 6 that the investment in MAX5 is outweighed by the number of reports that can be generated over that time frame.

**Tests for web-based application throughput and response time**

In this test, our goal was to determine how much MAX5 improves the performance of an OLTP application that accesses a large database.

We used an OLTP application that used a WebSphere eXtreme Scale data grid as an inline cache, meaning that the application directed all read/write requests to the grid. All equipment was production-quality. If the data referenced in the request was not already in the grid, a special database loader brought the data in by accessing the database. This action incurs substantial processing overhead because the database frequently accesses a disk to obtain the data.

## Test setup and configuration

The server used in this test was a production-quality System x3690 X5 running Red Hat Enterprise Linux 6 and WebSphere eXtreme Scale 7.1. The DB2 system was on a separate IBM server with a storage system that supported 32 disks. Transactions were generated by a remote client system. We set up two memory configurations:

- ▶ With MAX5: 256 GB RAM internal to the System x3690 X5, plus 256 GB RAM installed in the MAX5 and attached to the System x3690 X5, for a total of 512 GB RAM
- ▶ Without MAX5: 256 GB RAM internal to the System x3690 X5 (called the *base system*)

This test configuration is shown in Figure 7. WXS represents WebSphere eXtreme Scale.

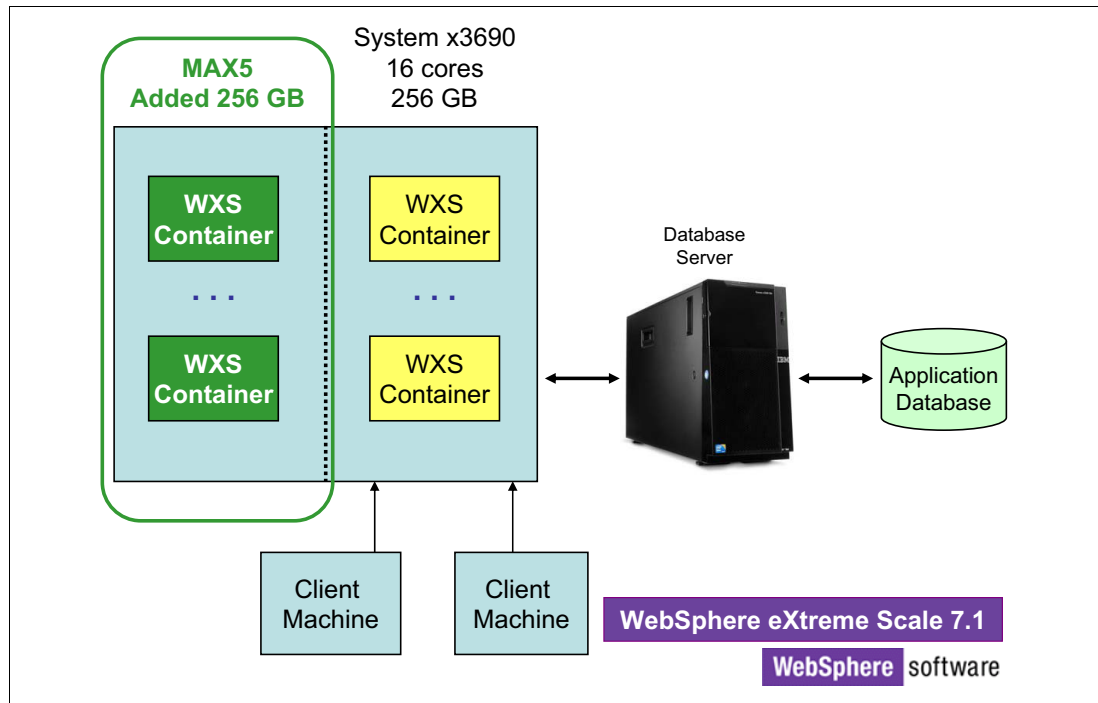


Figure 7 Test configuration that shows the System x3690 X5 with MAX5 and the 350 GB database

## Test day

The first part of the experiment was to run WebSphere eXtreme Scale on the System x3690 X5 with MAX5 attached, thus providing a system with 512 GB RAM. We created a WebSphere eXtreme Scale grid that uses 450 GB of RAM (45 containers with 10 GB each). The database was approximately 350 GB in size and fit completely in the grid. During the experiment, all data reads were satisfied from the grid.

The experiment contained two separate run scenarios, each with a different workload. In one test scenario, the workload consisted of 80% read requests and 20% write requests. On write requests, WebSphere eXtreme Scale first updated the grid and then synchronously wrote the data to the database. In the second scenario, the workload consisted of 100% reads.

In the second part of the experiment, the tests were repeated on the System x3690 X5 without the MAX5. This configuration reduced the amount of memory by one-half, leaving 256 GB in the system. We reduced the grid on that system by half as well, thus creating a grid that used 225 GB of RAM (45 containers with 5 GB each). There was not enough memory for the entire 350 GB database. As a result, some requests to the grid failed, and WebSphere eXtreme Scale read the data into the grid from the database. This state incurred the processing overhead of a remote call and possibly a disk access. To ensure that the grid

always had space available for newly requested data, we used the WebSphere eXtreme Scale built-in time-to-live evictor. The evictor removed any object in the grid not referenced for one hour.

## Results

Based on actual measurements, throughput with MAX5 was 2.4 - 4.9 times that of the system without MAX5. Response times were 60% - 80% less than the response times of the base system. Extrapolated results also showed that it can cost 49% more for a system without MAX5 to match the performance of a system with MAX5.

### **MAX5 increases throughput and reduces response time**

For the workload scenario of 80% reads and 20% writes, throughput on the MAX5 system was 2.4 times greater than the throughput of the base system, as shown in Figure 8. Response time on the MAX5 system was 60% less than on the base system, as shown in Figure 9.




Figure 8 Throughput: 80% read 350 GB database (larger is better)

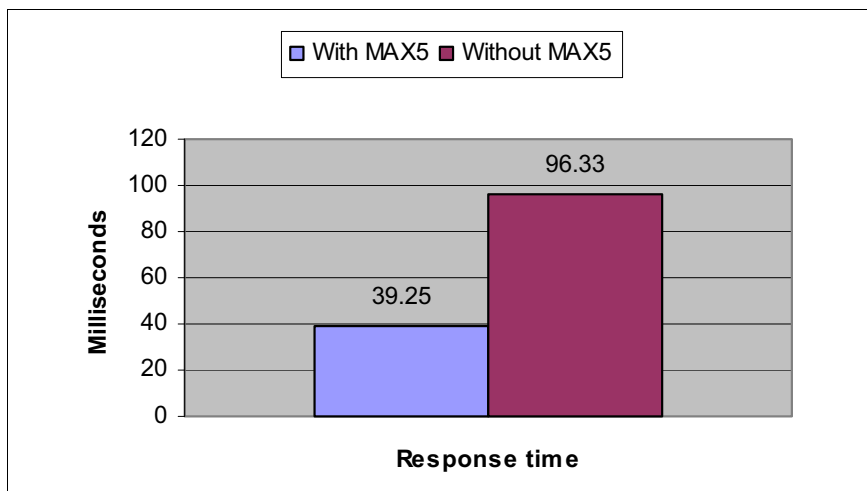


Figure 9 Response time: 80% read 350 GB database (smaller is better)

For the workload scenario of 100% reads, throughput on the MAX5 system was 4.9 times greater than the throughput of the base system, as shown in Figure 10. Response time on the MAX5 system was 80% less than the response time on the base system, as shown in Figure 11.

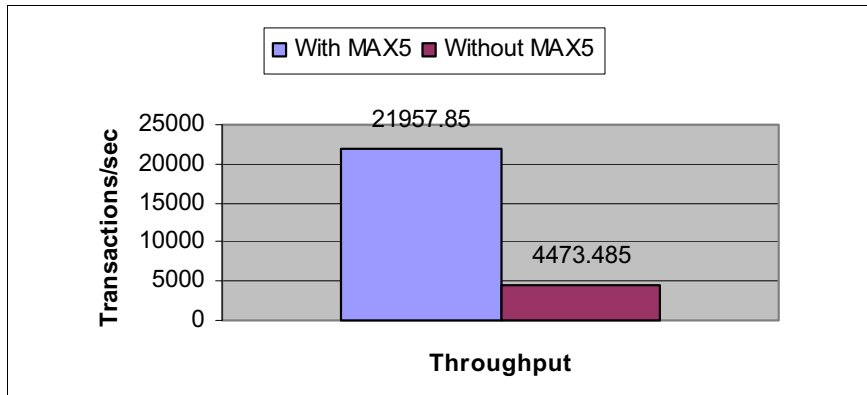


Figure 10 Throughput: 100% read 350 GB database (larger is better)

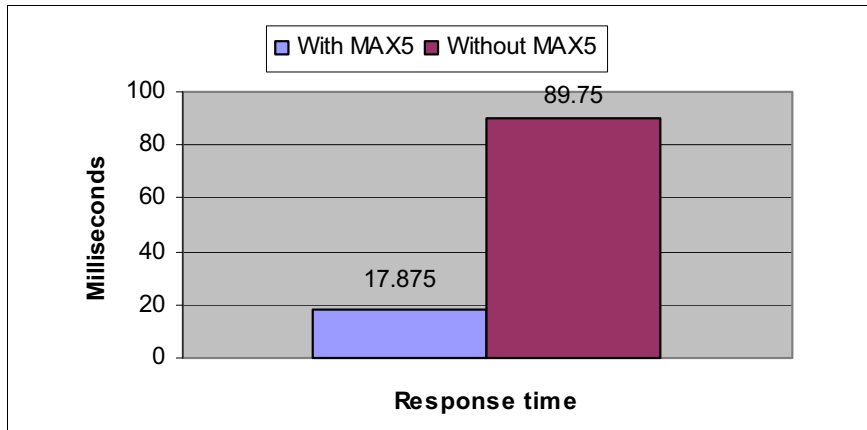


Figure 11 Response time: 100% read 350 GB database (smaller is better)

**MAX5 reduces the total cost of acquisition**

The STL team examined how much hardware had to be added to the base system to achieve the same 2.4 times gain in throughput observed in the MAX5 system for the 80% read test case. A comparable system without MAX5 is shown in Figure 12 on page 13.

The cost of this new system can be as much as 49% more than the cost of the MAX5 system that was used in our tests. Table 1 on page 13 shows our purchase requirement calculations. Values in the table are list prices that were published on the IBM website when the analysis was conducted in mid-2011. They might not be representative of current list prices. You can discuss your current environment with your IBM representative to determine your savings with MAX5 at the current prices. A large portion of the increased cost is from the purchase of additional WebSphere eXtreme Scale licensing. WXS represents WebSphere eXtreme Scale.

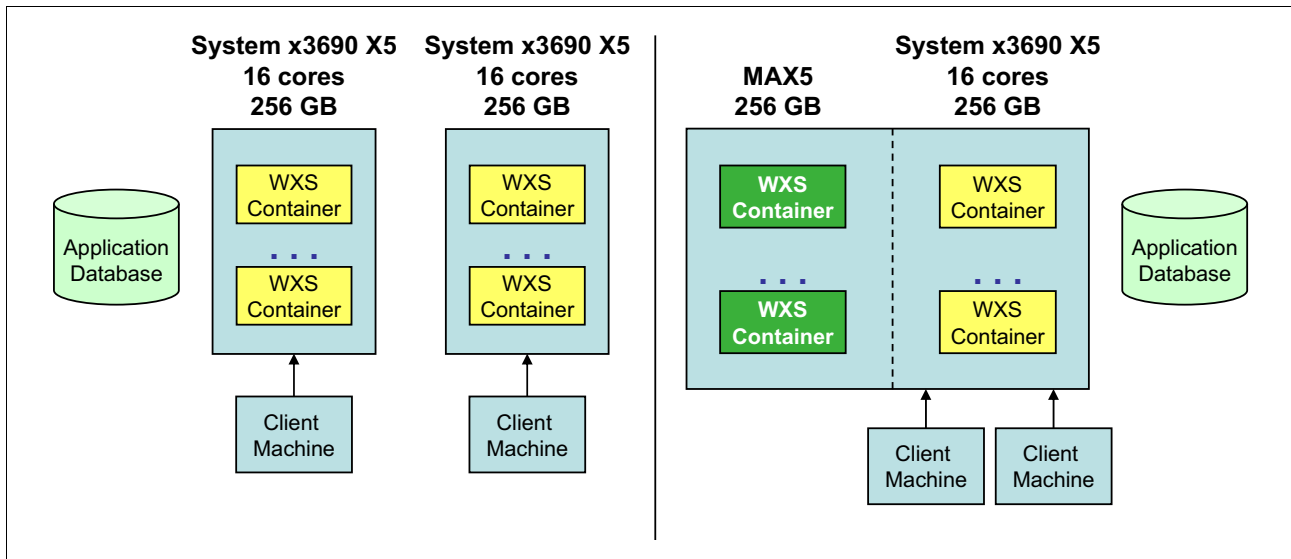


Figure 12 A system without MAX5 (left) needs to provide the same performance as the system with MAX5 (right)

Table 1 Details of purchase requirement calculations

Hardware/software	Cost	Quantity without MAX5	Total	Quantity with MAX5	Total
System x3690 X5 (includes 256 GB RAM)	\$17,364.00	2	\$34,728.00	1	\$17,364.00
MAX5	\$7,495.00	0	\$0.00	1	\$7,495.00
256 GB RAM	\$12,000.00	0	\$0.00	1	\$12,000.00
Disks			\$202,856.00		\$202,856.00
Red Hat Enterprise Linux	\$0.00	2	\$0.00	1	\$0.00
WebSphere eXtreme Scale	\$236,800.00	2	\$473,600.00	1	\$236,800.00
Total			\$711,184.00 <sup>a</sup>		\$476,515.00 <sup>a</sup>

a. Values in the table are list prices that were published on the IBM website when the analysis was conducted in mid-2011. They might not be representative of current list prices. You can discuss your current environment with your IBM representative to determine how much MAX5 can save you at the current prices.

## Summary

Our tests demonstrated that MAX5 can offer significant benefits to your organization:

- ▶ Response time can be decreased for business analytics and data warehouse workloads by the addition of MAX5. In our tests, the ability to cache all of the data in memory cut the response time from 1.5 - 2.8 times.
- ▶ Throughput can be increased by using MAX5 for business intelligence workloads. In our week-long test, we produced twice as many reports with MAX5 than without. Therefore, over a three-year period, the cost of reports can be 31% - 64% lower with the addition of MAX5.
- ▶ MAX5 can provide significant performance gains for transaction processing applications with large databases. With MAX5, throughput was 2.4 - 4.9 times greater than that of the test system without MAX5. Similarly, MAX5 enabled response times that were 60% - 80% less than the response times on the system without MAX5.
- ▶ MAX5 technology can reduce your total cost of acquisition because it adds memory only, rather than additional hardware, software, and the associated operating costs.
- ▶ MAX5 can increase performance at a lower cost of acquisition to meet your business needs. Our tests showed that it can cost 49% more for a system without MAX5 to match the performance of a system with MAX5.

All testing was completed by using production-quality equipment. Test results are based on actual measurements. Estimations of future results are based on extrapolations from actual measurements. MAX5 technology is one of the features of the eX5 product line that puts it in a class by itself. It is an optional unit for the IBM eX5 line of servers.

## For more information

For more information about the products referenced in this publication, visit these websites:

- ▶ The IBM Sales team:  
<http://ibm.com/systems/x/>
- ▶ IBM MAX5:
  - *The MAX5 Advantage: How IBM System x MAX5 Benefits Microsoft SQL Server Data Warehouse Workloads*, REDP-4751:  
<http://www.redbooks.ibm.com/abstracts/redp4751.html>
  - *IBM Unveils Industry's First Systems that Rewrite Economics of 'Industry-Standard' Computing*, 2010:  
<http://www.ibm.com/press/us/en/pressrelease/29570.wss>
- ▶ IBM System x3690 X5 server:
  - *IBM System x3690 X5: Ground-breaking two-socket server delivers outstanding performance, memory and storage*:  
<http://www.ibm.com/systems/x/hardware/enterprise/x3690x5/>
  - *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and BladeCenter HX5*, REDP-4650:  
<http://www.redbooks.ibm.com/abstracts/redp4650.html>

- ▶ IBM System x servers:
  - *High density, high performance: IBM System x white paper:*  
<http://ibm.com/systems/x/>
- ▶ IBM eX5 Enterprise Systems:
  - *IBM eX5 enterprise systems: Breakthrough innovation for smarter computing:*  
<http://www.ibm.com/systems/info/x86servers/ex5/>
- ▶ IBM Systems home page:
  - *IBM servers, storage, and solutions for your business:*  
<http://www.ibm.com/systems/>
- ▶ WebSphere eXtreme Scale:
  - Product information page:  
<http://www.ibm.com/software/webservers/appserv/extremescale/>
- ▶ DB2:
  - Product information page:  
<http://www.ibm.com/software/data/db2/linux-unix-windows/>

## The team who wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.



**Tim Bohn** is an IBM Senior Certified IT Specialist and Open Group Master Certified IT Specialist with the IBM Software Group, Strategy and Technology Lab (STL). His current focus is the IBM System x platform hardware and software. Before joining STL in 2010, Tim was the Global Community Practice Leader for the IBM Rational® Systems Development Community, where he helped clients with systems and software development practices. Prior to joining IBM Rational Software, Tim worked as a software and systems engineer for 15 years. Tim holds a Bachelor's and Master's degree in Computer Science from USC. He is a yearly guest lecturer at USC and UCLA.



**Dr. Howard Operowsky** is a Senior Technical Staff Member in the IBM STL where he evaluates server and middleware technology and is currently focused on virtualization. Since joining IBM, he has been manager or technical leader of a variety of IBM products, including WebSphere Commerce architecture, collaborative multimedia applications, IBM AIX/ESA® operating system architecture, the design and development of compilers and programming languages for high-end parallel systems, and performance analysis of mainframe operating systems. He has co-authored multiple patents in the area of multimedia applications and is the recipient of an IBM Outstanding Technical Achievement Award for work on operating system architecture. He received his Ph.D. in Computer Science from the Courant Institute of Mathematical Sciences at New York University. He is a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE).



**Karen Lawrence** is a Technical Writer with the IBM ITSO team in North Carolina. She has 25 years of experience in IT, with expertise in application design, change management, and the software development lifecycle (SDLC). She has worked with SMEs on leading technologies in global, regulated enterprises in the areas of storage systems, security, and disaster recovery, and the applications to support these areas. Karen holds a Bachelor's degree in Business Administration from Centenary College in New Jersey.

Thanks to the following people for their contributions to this project:

Many thanks to Randall W. (Randy) Lundin, Product Marketing - High End System x, for his contributions and sponsorship of this project.

Thanks also to Linda Robinson, IBM Redbooks® graphics specialist, and to the editing team from the International Technical Support Organization for their contributions to this project.

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

<http://www.ibm.com/redbooks/residencies.html>



## Stay connected to IBM Redbooks

- ▶ Find us on Facebook:  
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:  
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:  
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<http://www.redbooks.ibm.com/rss.html>



# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2012. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-4846-00 was created or updated on July 18, 2012.



Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review Redbooks form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an email to:  
[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400 U.S.A.



## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX/ESA®	Rational®	System x®
Cognos®	Redbooks®	WebSphere®
DB2®	Redpaper™	X-Architecture®
IBM®	Redbooks (logo)  ®	z/OS®

The following terms are trademarks of other companies:

Intel Xeon, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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