



Tim Bohn
Karen Lawrence

Advantages of IBM eX5 for Database Workloads

Executive summary

Methods for optimizing database workloads can vary, depending on the type of database and the purpose of the workload. For online transaction processing (OLTP) workloads, for example, there are thousands of I/Os per second (IOPS). We describe ways to optimize response time to the user for this type of workload. From the Business Analytics (BA) workload perspective, sometimes you might need to run multiple, sometimes complex, queries that can take a long time, even hours. For this type of workload, optimized I/O throughput is the goal.

IT departments cannot acquire the hardware and software to meet all needs for all projects. To address this requirement, select the equipment that meets as many project needs as possible, realizing that there are trade-offs. For database workloads, a delicate balance exists between CPU processing power and I/O throughput, so the trade-off can be considerable.

When analyzing the difference in OLTP and BA workloads, you must realize that it is impossible to have the best processing power and the highest IOPS from the same hardware and software. The tuning for these workloads differs. To compare, one component of OLTP workloads is database locking. Efficient locking is required because updates are going against the data, and inefficient locking degrades as the number of users grows.

In contrast, for BA, query optimization becomes much more important. Tools, such as Materialized Query Tables and Multidimensional Clustering Tables, help to speed up complex queries where the data might not change (is precomputed) but tend to be larger (just like the data in BA). Another factor, which becomes more prevalent, is query monitoring and management. Proper monitoring ensures that queries use only a certain amount of resources before either being stopped or moved down to a lower priority to limit the impact to the system. Additional effort is spent on both the denormalization and partitioning, to try and limit I/O effects.

This IBM® Redpaper™ publication describes IBM products that address requirements specific to OLTP and BA workloads. These products are the result of experience with clients and IBM Business Partners over many years. They convey the importance of database workload efficiency in multiple workload environments.

This paper details the features of IBM eX5 servers. IBM designed these servers to increase efficiency and lower operating costs, therefore, maximizing workload optimization for OLTP and BA workloads. Additionally, this paper describes other solutions that work in concert with the eX5 family of servers. This paper is based on in-house studies conducted by the IBM Strategy and Testing Laboratory (STL). For these studies, production-quality IBM System x3690 X5 servers were used.

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

The balanced system design objective

For database systems, a delicate balance exists between CPU processing power and the I/O throughput needed from the disk. This balance involves other factors, such as memory. When you add more CPU processing power, you become I/O limited. The processor spends more time waiting for data from the I/O system. As you add more processing power to your I/O system, the system becomes CPU starved.

In a balanced system design¹, the most efficient, optimized systems are not starved for I/O or CPU. The systems are considered balanced so that the CPU and the storage do not become a bottleneck. Figure 1 describes this concept.

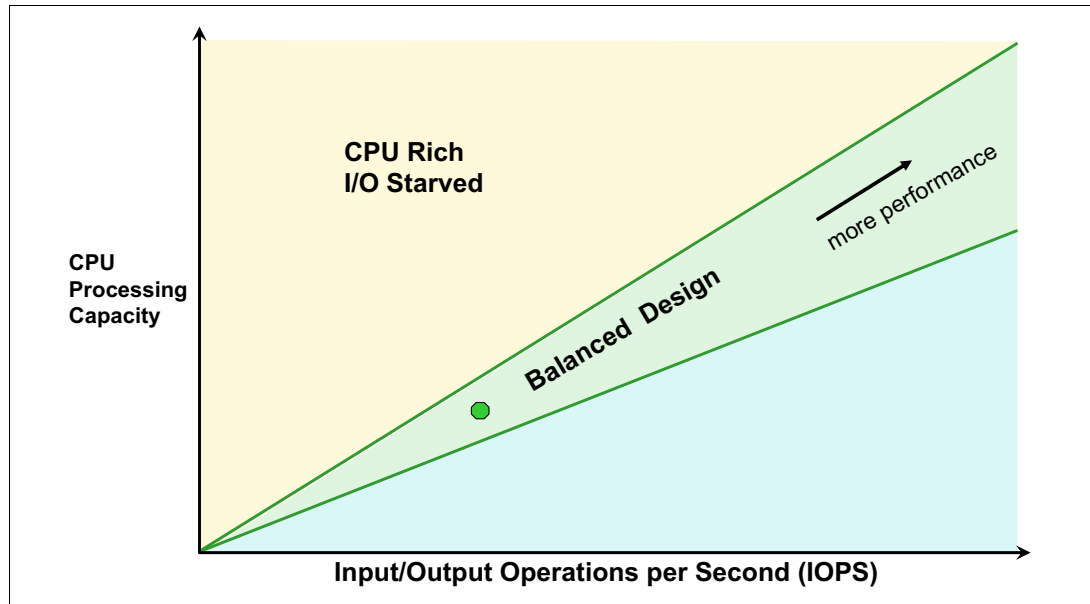


Figure 1 Efficient database OLTP performance requires a balance of optimizations

¹ For more information about balanced system design, see *The Benefits of Optimizing OLTP Databases Using IBM eXFlash Solid-State Drives*, REDP-4849, April 2012, available at this website: <http://www.redbooks.ibm.com/abstracts/redp4849.html>

Characteristics of data processing workloads

Each type of workload includes certain characteristics or tasks that define the workload. Figure 2 lists characteristics that are applicable to OLTP and BA workloads. Notice the heavy use of system resources by each type of workload.

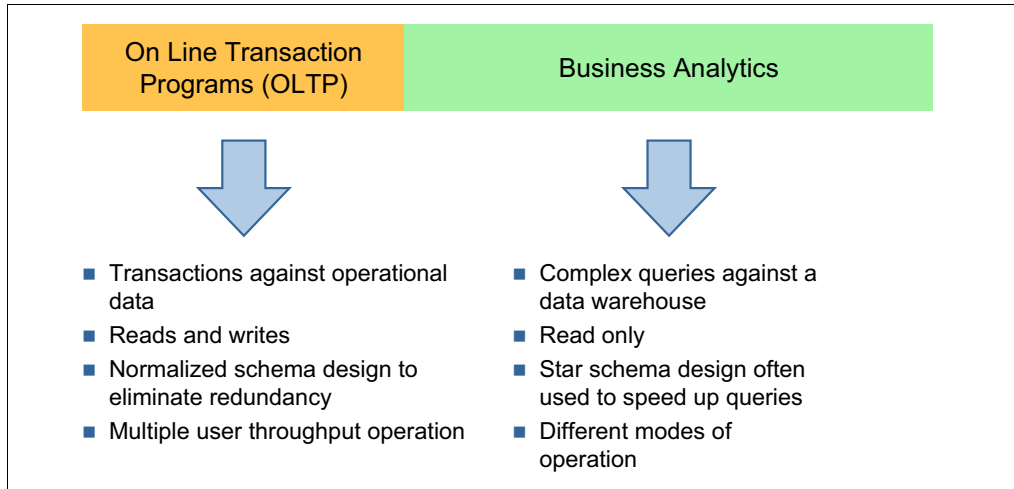


Figure 2 Data processing workloads have different characteristics

We examine the details of the tasks to help you fully understand how each characteristic is an important factor in workload optimization.

Characteristics of OLTP workloads

OLTP workloads are characterized by many transactions from many parallel users, with those transactions that go against the data. A mixture of reads and writes exists. The database needs to be normalized to eliminate data redundancy and to update efficiently.

Characteristics of BA workloads

Three possible methods for running BA workloads are defined as serial execution, concurrent execution, and fixed execution. The following definitions describe these types of workloads:

- ▶ *Serial execution:* A single user uses the system to run successive reports.
- ▶ *Concurrent execution:* Multiple users run reports in parallel; you want to know how many reports per hour can be generated.
- ▶ *Fixed execution:* Similar to concurrent execution, but multiple users run reports in parallel; you want to know how quickly you can issue a fixed number of reports.

Optimizing database workloads for the best performance and price

OLTP and BA workloads have different, often conflicting, characteristics. Therefore, what is the best way to optimize your systems to supply the best performance for these workloads and at the best price? The IBM STL group tested both the OLTP and BA scenarios to determine how the need for the best performance at the best price can be met.

Optimizing OLTP workloads

For OLTP workloads, we conducted a study that reviewed the Transaction Processing Performance Council (TPC-C)² standard benchmarks for systems that are considered balanced in terms of CPU processing power and IOPS. This balance means that neither system is a bottleneck.

We looked for the system setups that provide the most transactions per CPU and hard disk drives used. We used the following calculations in our testing configurations:

1. We took the better half (top 50%) of the systems with the largest transaction per CPU Performance (CP) rate. We then calculated the CPU better half transactions per minute (tpm) C/CP average.
2. We took the better half (top 50%) of the systems with the largest IOPS rate. We then calculated the hard disk drive better half tpm C/HD average.
3. We found the top 30% of balanced systems with the tpm C/CP and tpm C/HD rates closest to the better half averages.

Figure 3 shows our results, where CP units are a standardized set of common performance units, created with multiple benchmarks across the different processors.

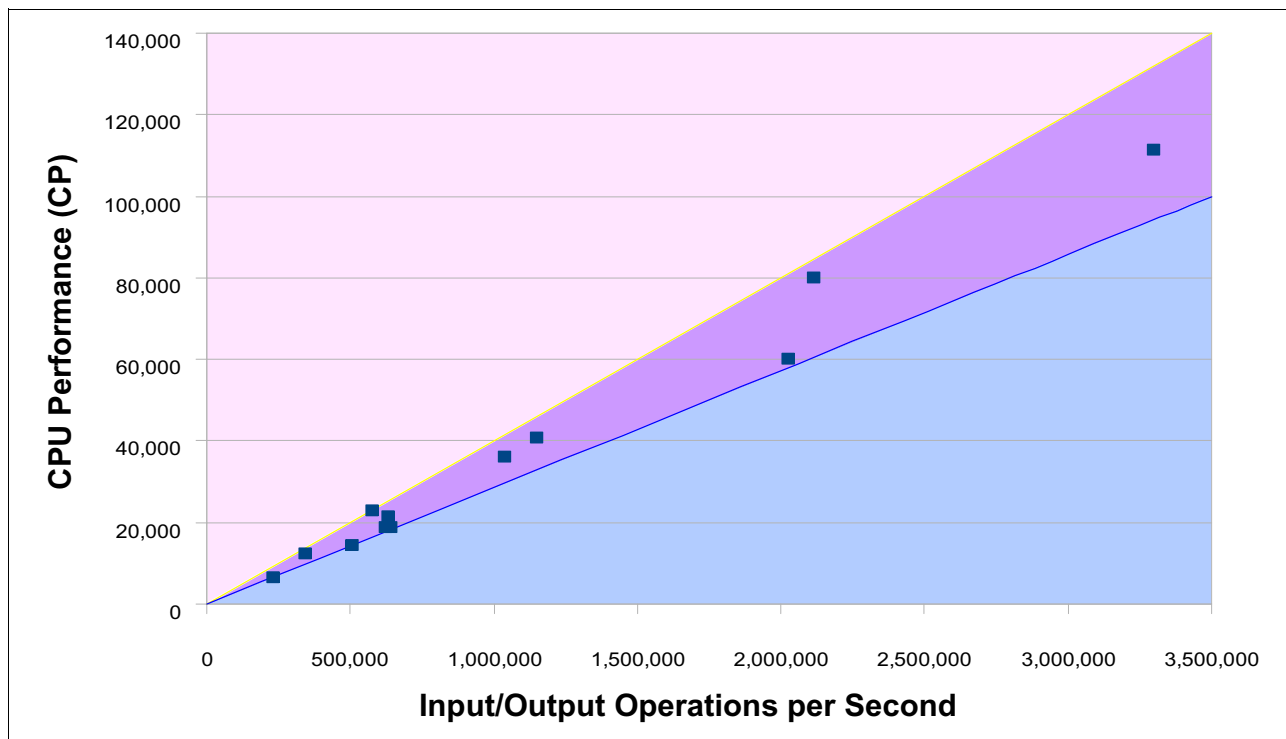


Figure 3 Best-optimized, hard disk drive (HDD) results for OLTP

² The TPC-C is a standard testing tool used in the industry. This tool simulates the “execution of multiple transaction types that span a breadth of complexity.” For more information, see this website: <http://www.tpc.org/default.asp>

Optimizing BA workloads

BA workloads have different modes of operation from which to compare performance. For BA workloads, the IBM STL group tested the I/O throughput made available in three modes of operation. The following scenarios were used to test the I/O throughput:

- ▶ *Serial execution:* A single user ran each of nine reports, consecutively, using 100% of the system resources. Serial execution shows how well a system issues the BA reports in isolation.
- ▶ *Concurrent execution:* Eighty users ran reports concurrently. We calculated an average number of reports completed per hour. We tested a concurrent mix of three different types of reports that were categorized as simple, intermediate, and complex. Each user repeatedly ran a particular report as many times as possible within a specified time. At the end of the allotted time, no more reports were started. When the running queries finished, the numbers for each type of report were tallied and the averages calculated.
- ▶ *Fixed execution:* A fixed number of reports were issued. Eighty users ran reports in parallel until all were finished. The total elapsed time was then measured.

Results of this test are described further in “Workload optimization: BA” on page 14.

Workload optimization: OLTP

In this section, we describe IBM hardware that was designed for use in workload optimization. We also cover specific technologies, in-house testing conducted by the STL group, and how tests results can be useful for your OLTP and BA workloads.

IBM System x and IBM DB2

Together, IBM System x® servers and IBM DB2® are designed to increase performance and efficiency. They can provide what you need to maintain a balanced system as your IT needs grow and change.

There are powerful symmetric multi-threading (SMT) systems that can scale up to the most cores and threads. Or, if the systems are used in conjunction with DB2 IBM PureScale, they can scale out with almost linear scaling, when staying transparent.

If your system needs more I/O throughput, eX5 servers can add eXFlash drives to boost onboard throughput. The IBM Storwize® V7000 can also be added for even more IOPS. Figure 4 on page 6 shows the performance and efficiency data for these two products. The figure also shows how CP units are a standardized set of common performance units, created with multiple benchmarks across the different processors.

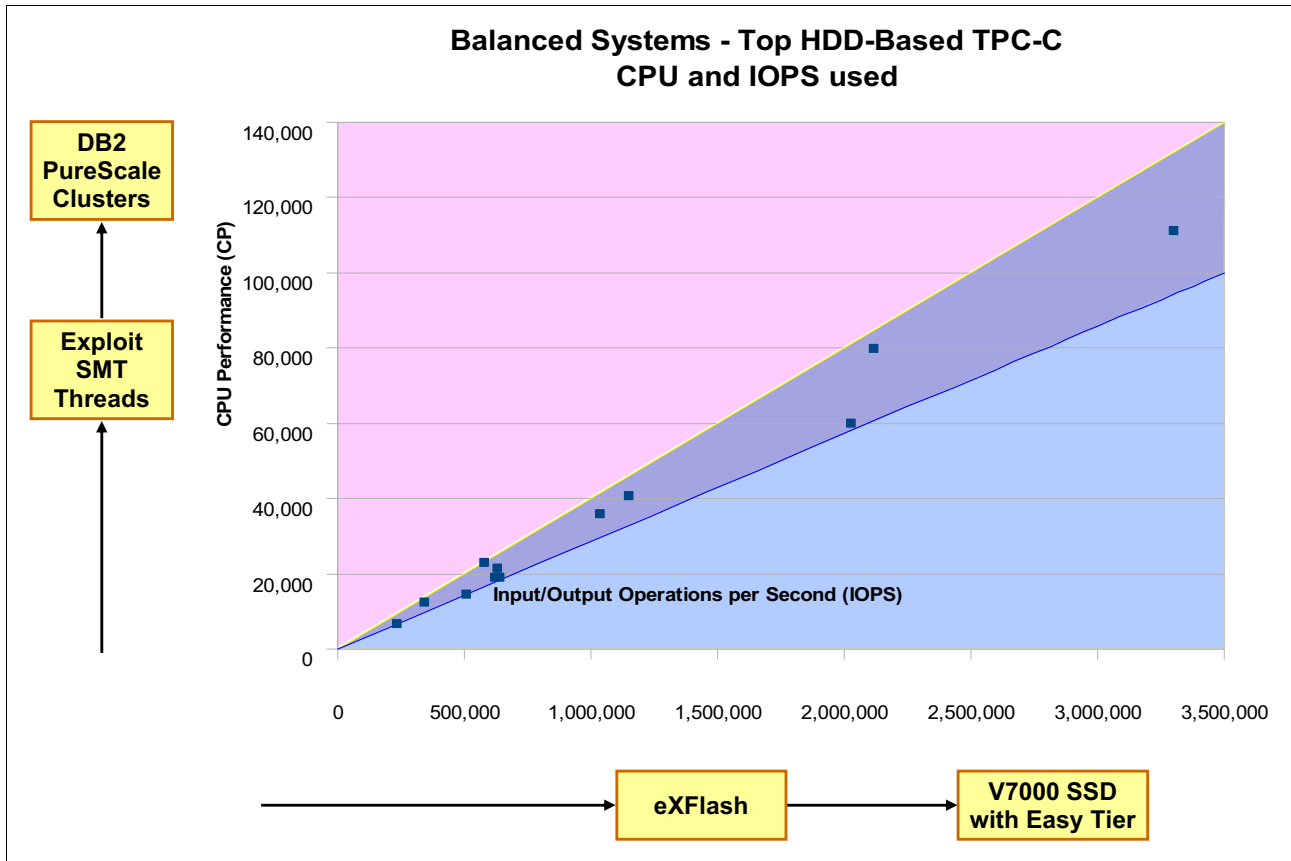


Figure 4 System x, eXFlash, and Storwize V7000 are designed to improve efficiency

System x with DB2 or eXFlash

The IBM eX5 has the CPU processing power to meet the performance requirements of OLTP workloads. The eX5 servers cover a wide range of needs, from a two-socket, 20-core blade and rack mount server; to a four-socket, 40-core System x3850 X5 that can scale to eight sockets and 80 cores. DB2 helps by scaling efficiently to enable the use of the largest System x server (in terms of processing power). Figure 5 on page 7 outlines more about the processing power of System x servers.

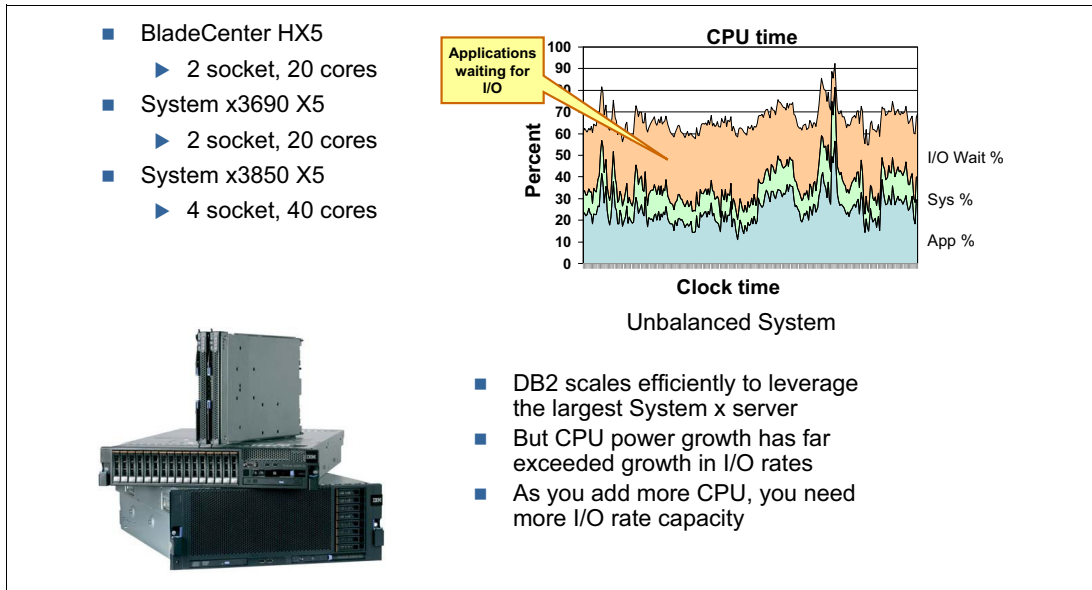


Figure 5 System x delivers processing power for your database workloads

As you add more CPU power, you realize that there is not enough I/O capacity. This lowered capacity becomes a bottleneck, leaving the CPU idle. We wanted to determine the performance impact of HDDs compared with eXFlash for OLTP workloads³. One solution to the lowered CPU idle time is to add eXFlash solid-state drives (SSDs). With eXFlash technology, you can put two drives in the space of one hard disk drive. The System x3850 X5 holds up to 16 eXFlash SSDs, and the System x3690 X5 holds up to 24 eXFlash SSDs. Figure 6 shows the specifics of eXFlash with SSDs. More details about System x with eXFlash are listed in Figure 7 on page 8.

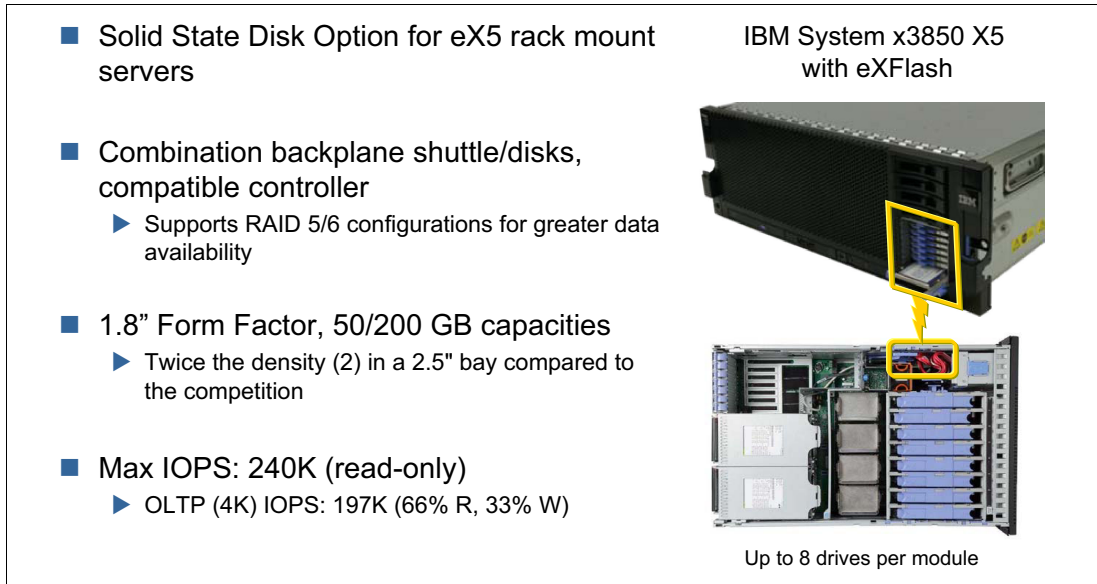


Figure 6 eXFlash adds SSD to System x

³ The study is briefly described in this paper. The details about the full study are available at this website: <http://www.redbooks.ibm.com/abstracts/redp4849.html>

For the database server, we used the System x3690 X5 with two Intel X7650 processors with eight cores for each processor. We compared the relative performance and cost of the following configurations:

- ▶ eXFlash SSD: 200 GB, Read, 30,000 IOPS per drive, eight drives per eXFlash (240K IOPS per eXFlash)
- ▶ HDD: 15K rpm serial-attached SCSI/Serial Advanced Technology Attachment (SAS/SATA) 3 Gbps link speed

Figure 7 further describes our test configuration.

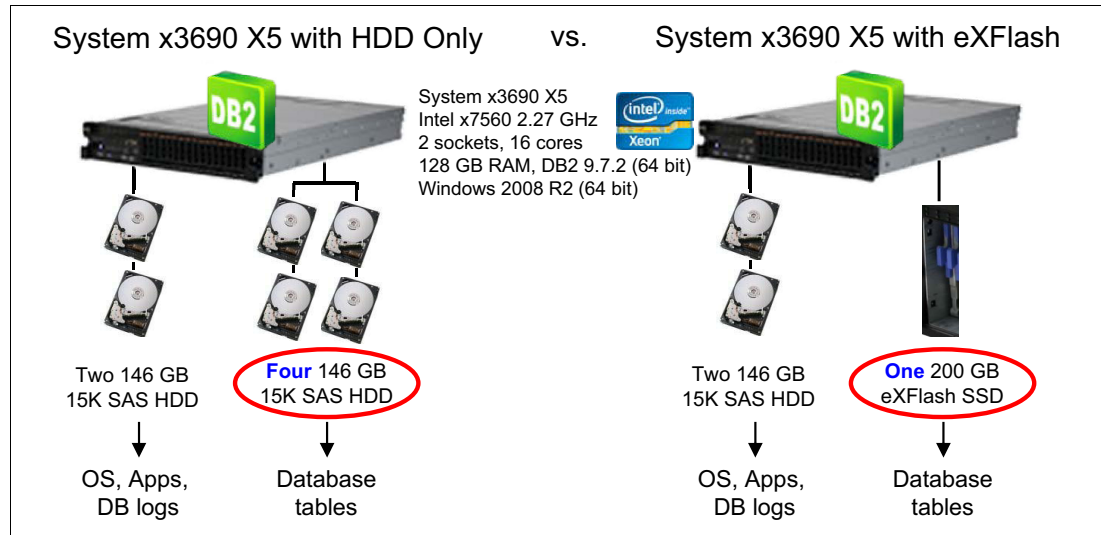


Figure 7 In this study, we assessed the performance of HDDs versus eXFlash for OLTP workloads

The testing focused on a 100 GB OLTP database to determine how many HDDs it needed to match the performance of a single eXFlash SSD. This scenario was entirely focused on the required performance levels. The extra storage capacity of the HDDs was not considered. Figure 8 on page 9 describes our calculations. DB2 can use eX5 with eXFlash to drive increases in performance.

The values in Figure 8 on page 9 indicate the list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was conducted in mid-2011. They might not be representative of current list prices. You need to discuss your current environment with your IBM representative to determine how much IBM can save you at the current prices.

Less than 1% cost increase, yields 4X performance improvement

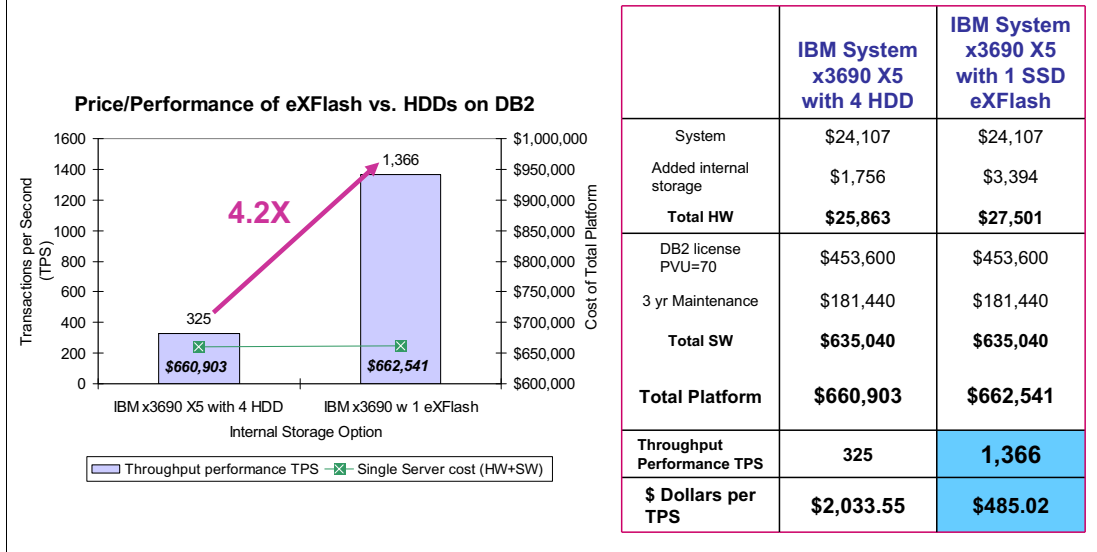


Figure 8 DB2 can use eX5 with eXFlash to drive higher performance at a lower cost

In another study, we compared IBM hardware and software with the similar hardware and software of a competitor. We compared the System x3850 X5, four socket (Intel X7560 2.26 Ghz, eight core), with SSDs and HDDs, to a four-socket comparison machine with an identical Intel processor.

The software in each configuration used the same OLTP driver, with DB2 on the IBM configuration and a comparable database on the comparison machine. Figure 9 shows more about the configuration details.

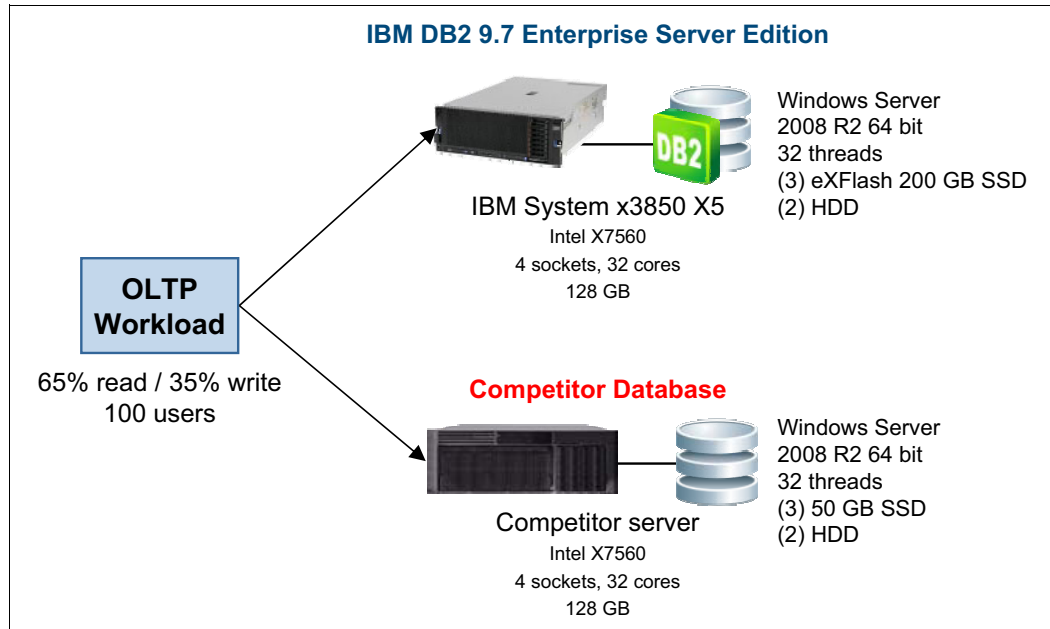


Figure 9 System x3690 X5 with eXFlash and DB2 compared with a similar system

Results shown in Figure 10 indicate that the IBM hardware and software stack outperformed the comparison stack by 54%. The hardware price difference was primarily because the IBM eXFlash SSDs were more expensive than the competitor SSDs. The IBM eXFlash SSDs also outperformed the SSDs of the comparison machine. Figure 10 shows more results of this study. Values in Figure 10 indicate the list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was conducted in mid-2011. They might not be representative of the current list prices. You need to discuss your current environment with your IBM representative to determine how much IBM can save you at the current prices.

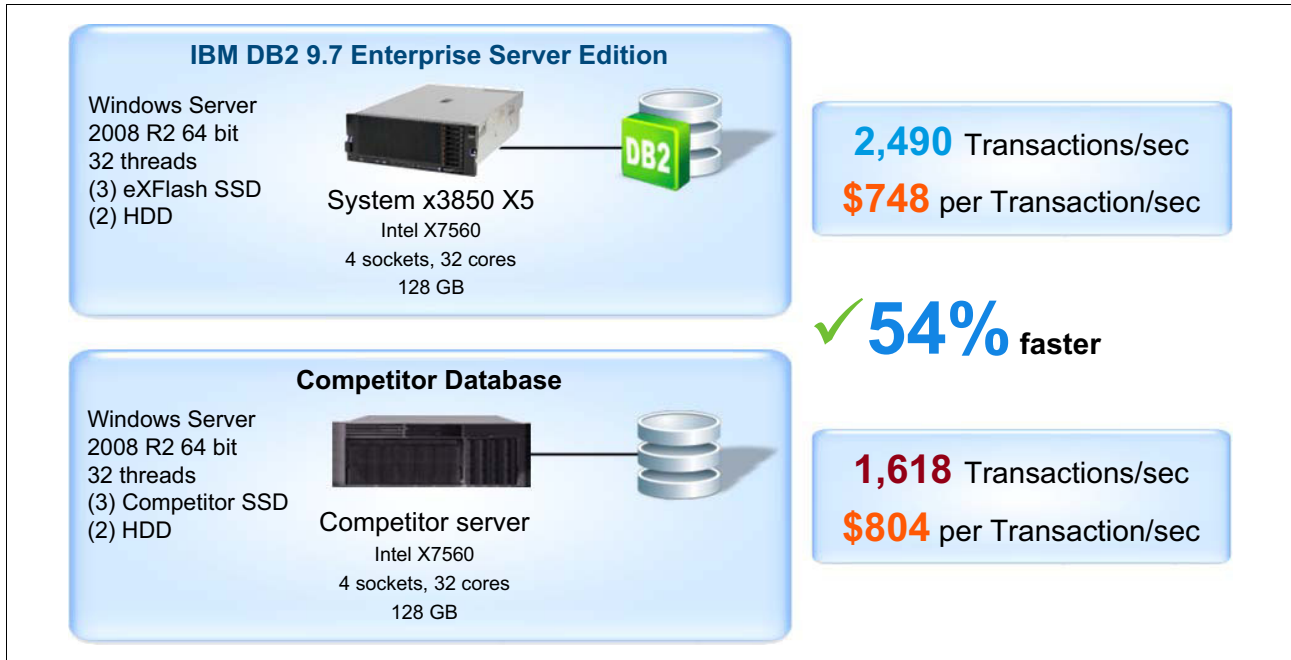


Figure 10 DB2 on System x increases performance by 54% over a comparable system

IBM DB2 Enterprise Edition on System x

In our test, we compared the cost of the DB2 Enterprise Edition to the equivalent comparison software. Results show that the normalized cost of the transactions per second (TPS) of the DB2 Enterprise Edition came out slightly ahead. See Table 1 for the comparison.

List prices: Values in Table 1 and Table 2 on page 11 are the list prices that were published on the website when the analysis was conducted in mid-2011. They might not be representative of current list prices. You need to discuss your current environment with an IBM representative to determine how much IBM can save you at the current prices.

Table 1 IBM DB2 Enterprise Edition and the comparison software

| | IBM DB2 Enterprise Edition | Comparison software |
|---------------------------------|----------------------------|--|
| Price for each PVU ^a | \$405 ^b | .5 (Core Factor Table) x 32 Cores: \$47,500 ^c |
| 100 PVUs (System x3850 X5) | \$40,500 | \$760,000 |
| Cores | 32 | 32 |
| Support, year 1 | Included | 22% of original cost of each year: \$167,200 |
| Year one total price | \$1,296,000 | \$927,200 |
| Support, years 2 and 3 | \$518,400 | 22% of original cost for each year: \$167,200 for each year or \$334,400 |
| Three-year total price | \$1,814,400 | \$1,261,600 |

a. Processor value unit

b. Values indicate the list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was performed in mid-2011. They might not be representative of current list prices. You need to discuss your current environment with an IBM representative to determine how much you can save at the current prices.

c. List prices and other information about products that are not IBM were obtained from a supplier of these products, published announcement materials, vendor web pages, or other publicly available sources. Information does not constitute an endorsement of such products by IBM. Questions on the capability of products that are not IBM need to be addressed to the supplier of those products.

By comparison, DB2 is typically sold as the Advanced Server Enterprise Edition. This version has all of the options included that clients normally need at a reduced cost, rather than adding the options separately. For the comparison software, several of those options are available only at an added cost. When this cost comparison was prepared, the IBM stack came out at 46% less cost per TPS, as shown in Table 2.

Table 2 Factor in the cost of frequently used features to extend the IBM total cost of ownership

| | IBM DB2 9.7 Advanced Server Enterprise Edition (AESE) | Competitor database |
|----------------------|---|-----------------------|
| Hardware | \$48,923 ^a | \$39,752 ^b |
| Core functionality | \$2,016,000 | \$1,261,600 |
| Advanced security | Included | \$305,440 |
| Compression | Included | \$305,440 |
| Partitioning | Included | \$305,440 |
| Disaster recovery | Included | \$265,600 |
| Total cost (3 years) | \$2,064,923 | \$2,483,272 |
| Total TPS | \$2,490 | 1,618 |
| Cost per TPS | \$829/tps | \$1,535/tps |

a. Values indicate the list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was performed in mid-2011. They might not be representative of current list prices. You need to discuss your current environment with an IBM representative to determine how much you can save at the current prices.

b. List prices and other information about products that are not IBM were obtained from a supplier of these products, published announcement materials, or other publicly available sources. Information does not constitute endorsements of such products by IBM. Questions on product capabilities that are not IBM need to be addressed to the supplier.

IBM Storwize V7000 and IBM Easy Tier

Next, we describe Storwize V7000 and IBM Easy Tier.

Enterprise-class workloads require SSDs with more storage capabilities

For larger storage needs, you can add the Storwize V7000. A Storwize V7000 expansion enclosure can hold twenty-four 6.35 cm (2.5 in.) drives. The drives can be a combination of HDDs and SSDs.

The control enclosure holds multiple expansion enclosures, can hold 240 drives per control enclosure, and is designed to move data from the HDD to the SSD. This configuration is based on real-time usage patterns that run in the background and move data automatically.

Figure 11 shows how enterprise-class workloads can benefit from an SSD that provides more storage capabilities.

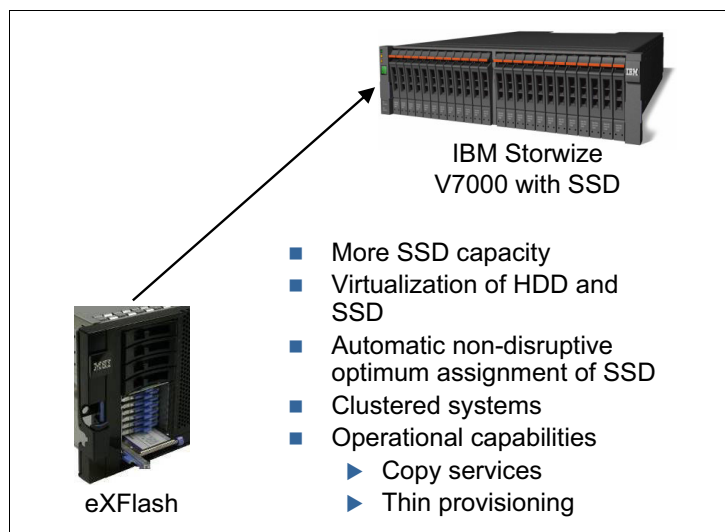


Figure 11 Enterprise-class workloads require SSDs with more storage capabilities

Storwize V7000 also has the thin provisioning feature, so that volumes appear to be larger than what is allocated. More space is allocated on demand for less wasted disk space. For more information about Storwize V7000, see this website:

http://www.ibm.com/systems/storage/disk/storwize_v7000/index.html

Easy Tier

Your SSD performance can be maximized further, while minimizing costs. Built into the storage controller, Easy Tier dynamically moves data to an SSD that is based on hotspots, and migrates data between storage devices. Performance improvements of up to 300% are observed by moving some of the data from HDD to SSD. Easy Tier is shown to improve response time by 78% by moving only 5% of the data from HDD to SSD. All storage vendors have SSD technology, but the Easy Tier is available from IBM only. Figure 12 on page 13 shows more of the optimization characteristics of Easy Tier.

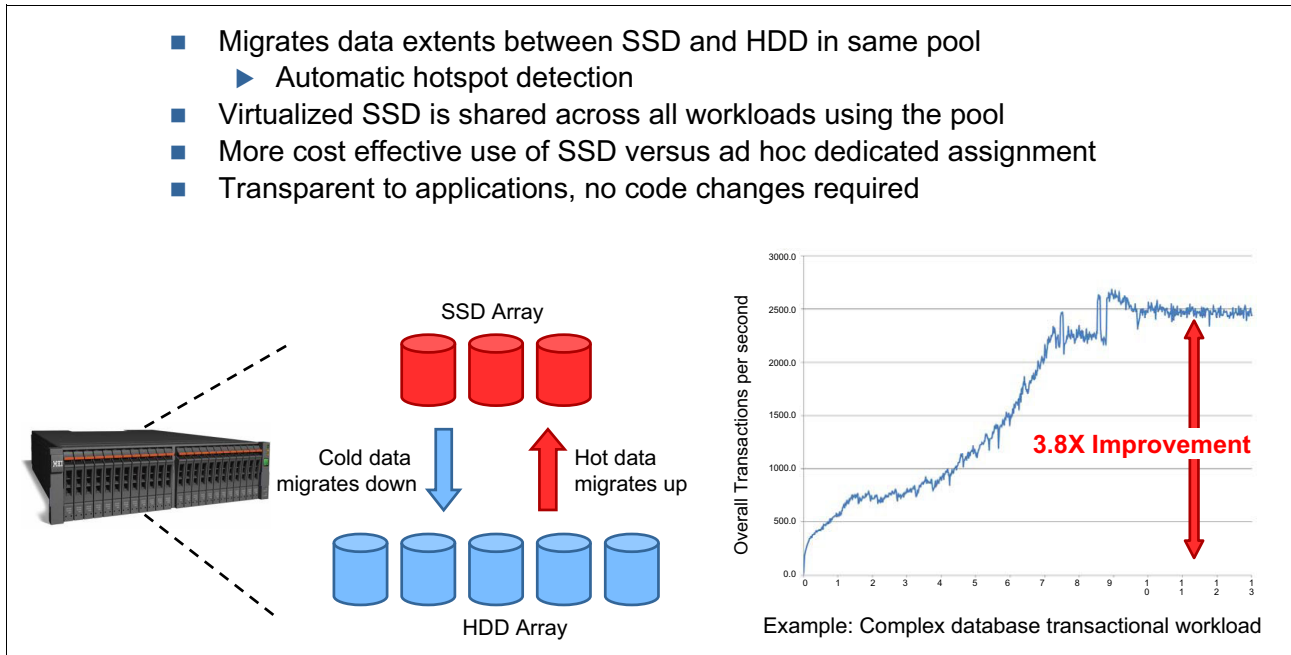


Figure 12 Easy Tier® in the Storwize V7000 automatically optimizes the use of SSDs

Figure 13 shows how SSD use becomes optimized over time, automatically. The left side of the figure shows the data that is in high demand and how it moves from HDD to SSD. The right side shows the TPS, which start at zero, and become optimized over time as a result of data movement to SSD.

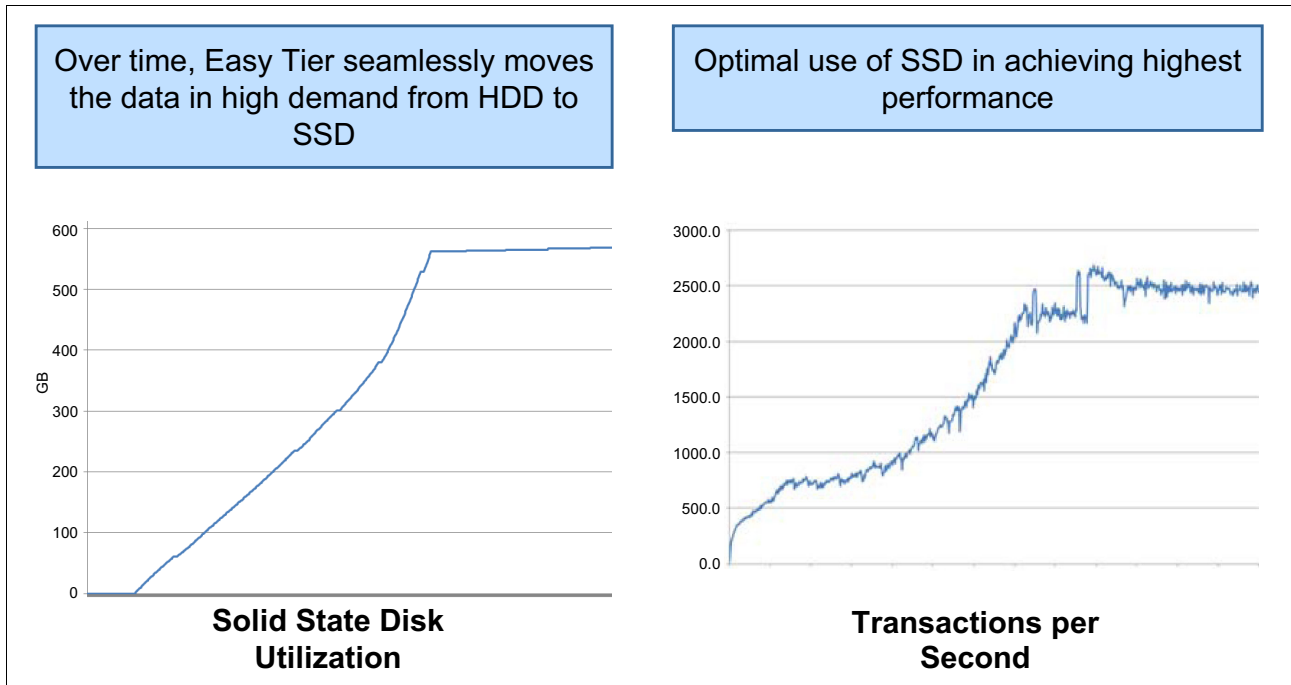


Figure 13 More results: Easy Tier in the Storwize V7000 automatically optimizes the use of SSDs

We compared performance by using an IBM BladeCenter® HX5 with a Storwize V7000 that was populated only with HDDs. This configuration was compared to the same setup, but we replaced four of the HDDs with SSDs. Figure 14 on page 14 reports these results.

Values in the figure show list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was conducted in mid-2011. They might not be representative of current list prices. You need to discuss your current environment with an IBM representative to determine how much IBM can save you at the current prices.

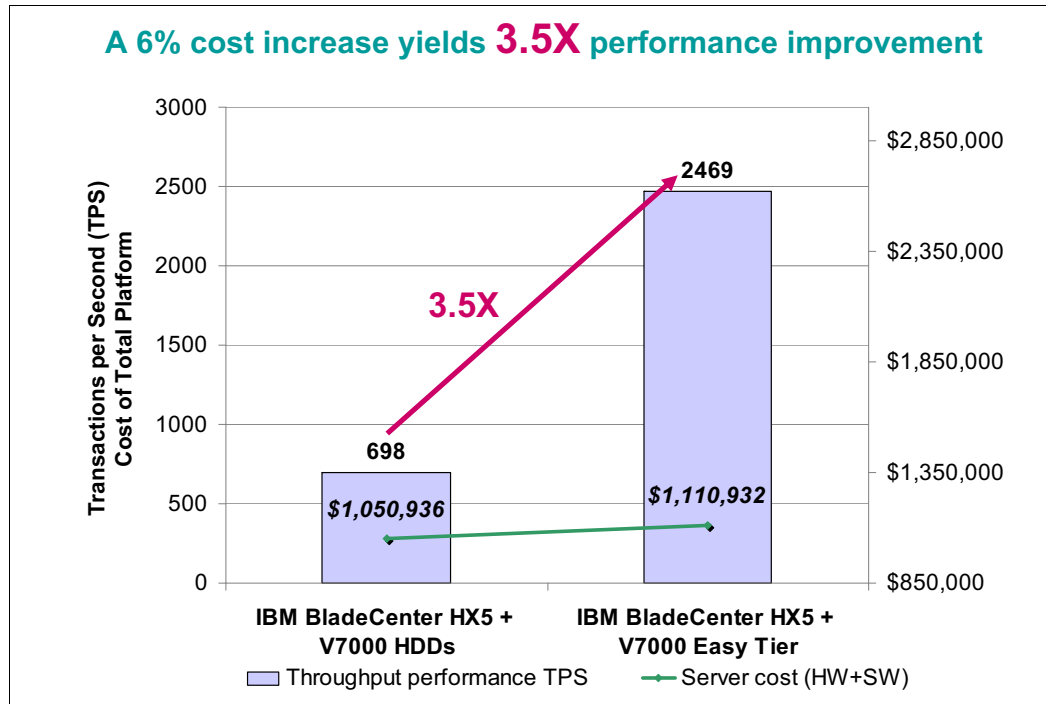


Figure 14 Storwize V7000 automatic data placement with Easy Tier is cost effective

The cost of this performance comparison increased by 6% because of the added SSDs. However, we also saw a 3.5-times boost in performance because of the SSDs. IBM Easy Tier is available at no extra charge.

Workload optimization: BA

In the study described in “Optimizing BA workloads” on page 5, we tested the I/O throughput that is available in the three modes of operation (serial, concurrent, and fixed execution). IBM offers different approaches to optimizing BA workloads, ranging from single-workload to multiple-workload systems. The following approaches to optimizing workloads are offered:

- ▶ *Optimized by design*, for example, DB2 on an eX5
- ▶ *Flexible integrated*, for example, an IBM Smart Analytics system
- ▶ *Tightly integrated appliances*, for example, IBM Netezza® 1000 or IBM DataPower®

These approaches are listed in decreasing order for flexibility, but in increasing order for time-to-value.

Flexible systems require more product knowledge and can be used for multiple workloads. The tightly integrated appliances require little product knowledge, because they are pre-configured to a specific workload.

The IBM STL group studied the process for optimizing systems for BA workloads, with a particular focus on key differentiation and how workloads provide economic value. Our philosophy is that one size does not fit all, because each project workload differs based on many factors. Figure 15 identifies some of these factors.

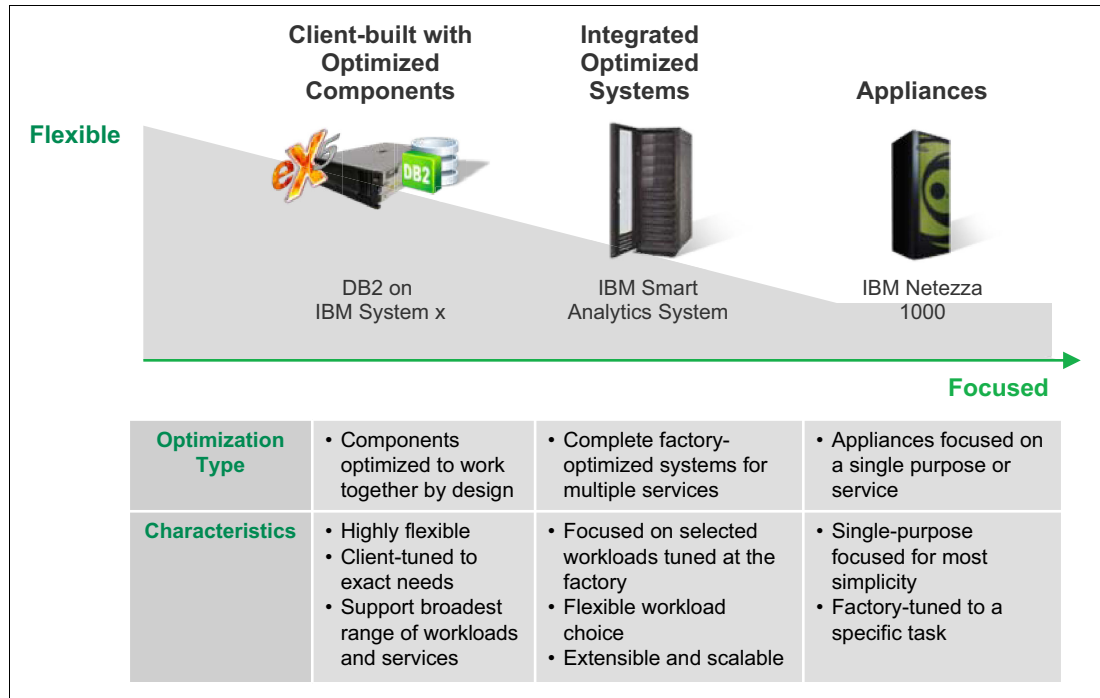


Figure 15 IBM has multiple BA solutions to meet your requirements

Several different configurations of optimized systems exist and focus particularly on key differentiations and how they can provide economic value.

The following options are provided with the IBM eX5 portfolio:

- ▶ *IBM System x3850 X5*: Versatile four-processor, 4U rack-optimized scalable enterprise server with optional memory expansion unit (MAX5) and workload-optimized models. This option provides a flexible platform to facilitate maximum utilization, reliability, and performance of compute-intensive and memory-intensive workloads.
- ▶ *IBM BladeCenter HX5*: Scalable blade server that enables standardization on the same platform for two-socket and four-socket server needs for faster time-to-value. This standardization occurs while delivering peak performance and productivity in high-density environments.
- ▶ *IBM System x3690 X5*: High-end, two-processor, 2U scalable server offers up to five times the memory capacity of the two-socket servers of today. This platform offers double the processing cores for superior performance and memory capacity.

The eX5 servers are highly scalable and flexible and provide the MAX5 option that allows for an additional 32 dual inline memory modules (DIMMs) of memory, decoupled from the processor. In addition, new eXFlash SSDs can be added to increase I/O throughput.

A flexible integrated system: IBM Smart Analytics System 5600

The key attributes of the IBM Smart Analytics System 5600 include the following features:

- ▶ Pre-packaged system
 - Includes optimal I/O subsystem, large I/O bandwidth, and modular storage scaling for analytics
- ▶ DB2 and IBM Cognos® capabilities
- ▶ A sound investment:
 - Enables you to start with the correct size and scale
 - Deploys in days, rather than months

A tightly integrated appliance: IBM Netezza 1000

Netezza appliances integrate database, server, and storage, delivering high performance and competitive pricing. Key attributes of the Netezza 1000 include the following benefits:

- ▶ Designed for rapid analysis of data scaling to petabytes.
- ▶ Delivers 10 - 100 times the performance improvements of traditional database vendors, at one-third the cost of options.
- ▶ Integrates business intelligence and advanced analytics.
- ▶ Delivers fast time-to-value; meaningful proof-of-concept can be available in as little as two weeks.
- ▶ Offers performance with its default settings, without requiring tuning, indexing, aggregations, and so on.
- ▶ Integrates with other products, such as Cognos and IBM Tivoli®; data assets are managed for their entire lifecycle.

Various BA workload solutions

IBM designed various solutions to optimize your BA workloads. This section reviews custom-built, integrated, and appliance solutions. Several systems on the market are designed to support the “one-size-fits-all” concept. In contrast, IBM offers custom-built, integrated, and appliance solutions. Figure 16 on page 17 defines these workload-optimized systems and associated products.

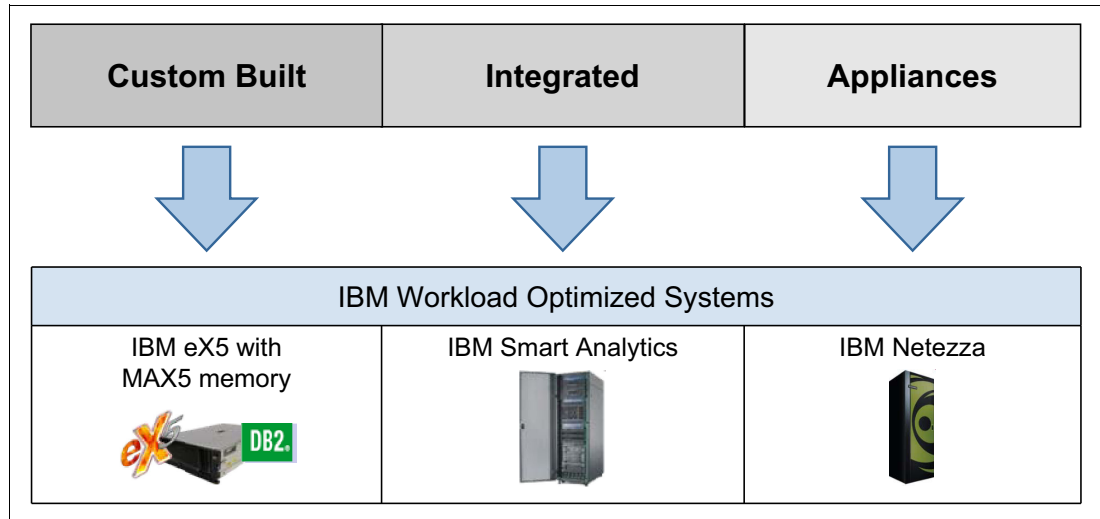


Figure 16 IBM offers custom built, integrated, and appliances solutions

Custom-built solution: MAX5

One of the primary features of the IBM eX5 family of servers is the memory expansion unit for the eX5, known as *MAX5*⁴. This optional unit, compatible with the eX5 family of servers, is designed to considerably increase server memory capacity. With MAX5, you can nearly double your memory capacity without adding processors.

The MAX5 provides the following features and benefits:

- ▶ Basic functions of MAX5:
 - Provides memory needed for database and virtualization workloads.
 - Allows higher memory capacity to be reached with less expensive DIMMs, providing more economical high-end implications.
- ▶ With MAX5, you can achieve these benefits:
 - Expand memory capacity (up to 1 TB of additional memory).
 - Up to double the number of memory DIMMs (32) than is possible with comparable machines. This increase equates to up to 512 GB of extra memory for the Intel Xeon 6500 and 7500 family of processors (by using 16 GB DIMMS). And it equates to up to 1 TB of extra memory for the Intel Xeon E7 family of processors (by using 32 GB DIMMS).
 - Add memory without adding processor cores, reducing software costs.
 - Increase database buffer size.

The STL group conducted a series of tests to determine the impact that adding memory, by using MAX5, had on performance of a BA workload. In this test, multiple, different types of reports were generated in succession. The performance and costs for the reports were calculated both before and after adding MAX5. For these tests, we used a production-quality IBM System x3690 X5 server and IBM DB2. These systems were used to test the effects of MAX5 on a 1.3 TB data warehouse in a BA environment. Figure 17 on page 18 depicts the process used.

⁴ For product specifications, see this website: <http://www.ibm.com/systems/x/options/memory/max5/specs.html>

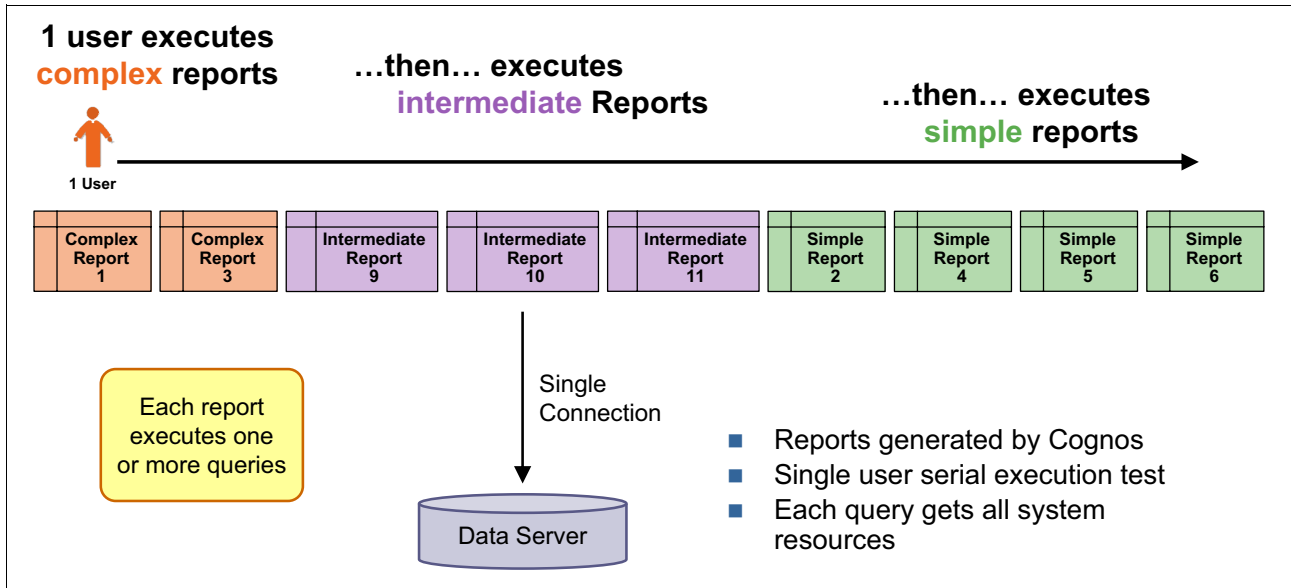


Figure 17 Process for testing serial execution of reports in isolation (single user)

As Figure 17 shows, this test was for a single user that requested various kinds of reports. The reports that were generated ran from complex, covering the entire data set; to simple, covering a small amount of data. The intent is to show how quickly the system responds with finished reports. The test was solely to test the effects of the added random access memory (RAM). The entire data set did not fit in the 512 GB RAM of the system without MAX5. When the RAM was doubled by adding MAX5, all of the data set was able to reside in the memory. This test is significant because MAX5 allows us to double the RAM, at a time when no other vendor offers this feature. Figure 18 describes the specifications.

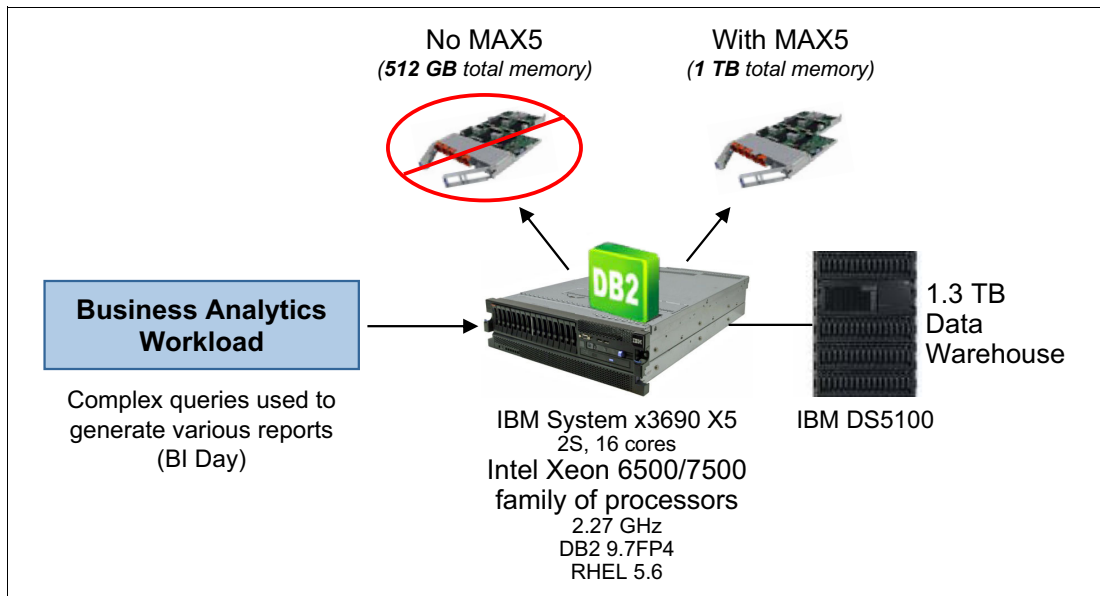


Figure 18 In this study, we assessed the impact MAX5 can have in BA workloads

With compression, a 1.3 TB data set nearly filled the buffer pool when using 1 TB RAM. However, the data set was larger than the buffer pool when using the base 512 GB RAM. Figure 19 on page 19 shows the results. In this example, lower is better.

As shown in Figure 19, with MAX5 attached, and when the response times for the three report types are averaged, reports were completed 1.5 to 2.8 times faster.

Results in Figure 19 are based on IBM internal tests of the System x3690 X5 with the MAX5 solution. This configuration is compared to the System x3690 X5 without MAX5, running an analytics workload in a controlled laboratory environment. The three-year total cost of acquisition includes expected hardware, software, service, and support, and is based on US list prices.

The cost calculation compares the average cost per report for the different report types, if the specified types of reports are run serially for the three-year life of the machine. The test compares the performance of running different reports of varying complexity. Results might not be typical and vary based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.

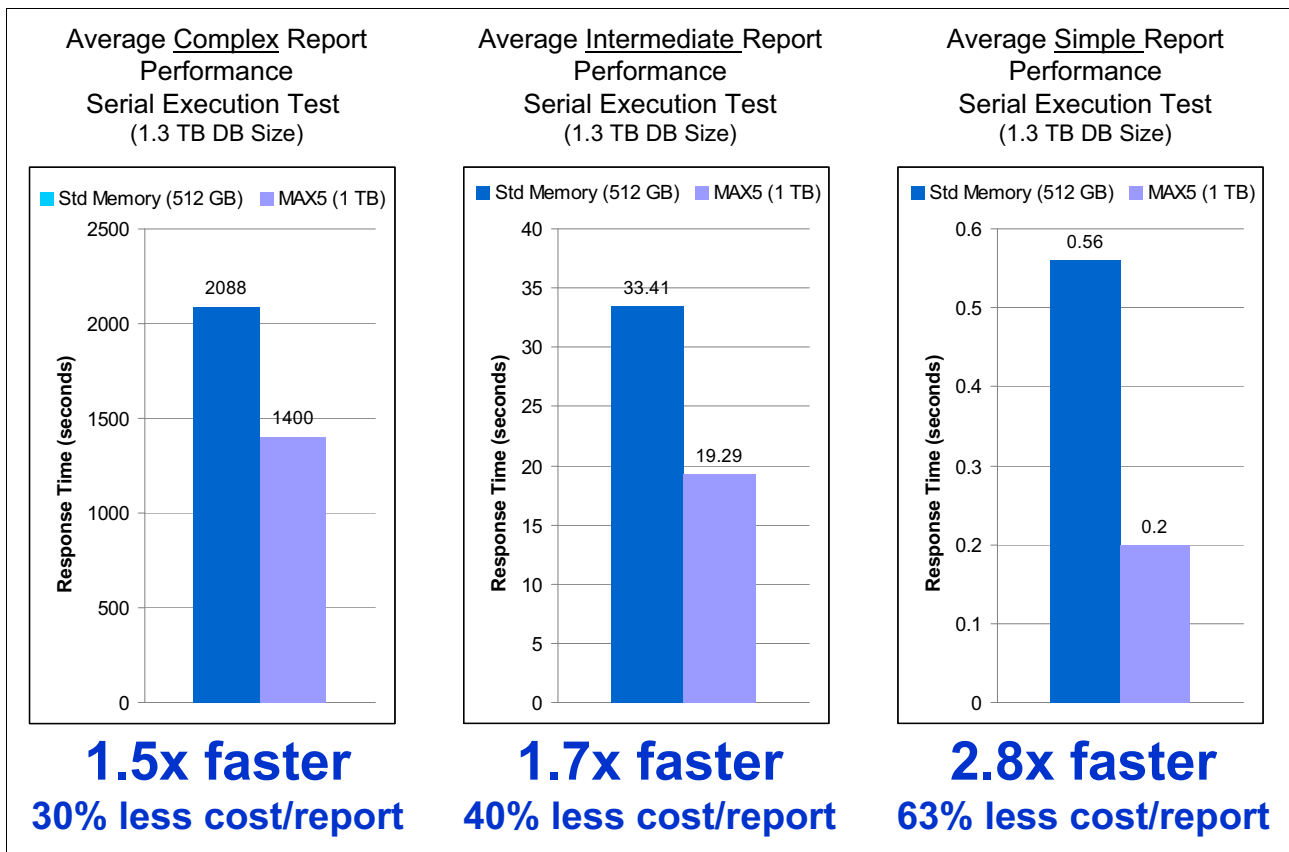


Figure 19 Comparison of report response times: MAX5 (1 TB) versus standard memory (512 GB)

The STL group ran the same test using the System x3690 X5 with MAX5 and comparison hardware and software. Figure 20 on page 20 shows the suggested configurations for these machines. Results in Figure 20 are based on a three-year total cost of acquisition (the competitor configuration is a previous version that is no longer available). Total cost also includes the hardware, software, service, and support that are used to run OLTP workloads. Calculations were derived in January 2012, based on US list prices during that time, from <http://www.ibm.com>. Prices vary by country. Users of this document need to verify the applicable data for their specific environment.

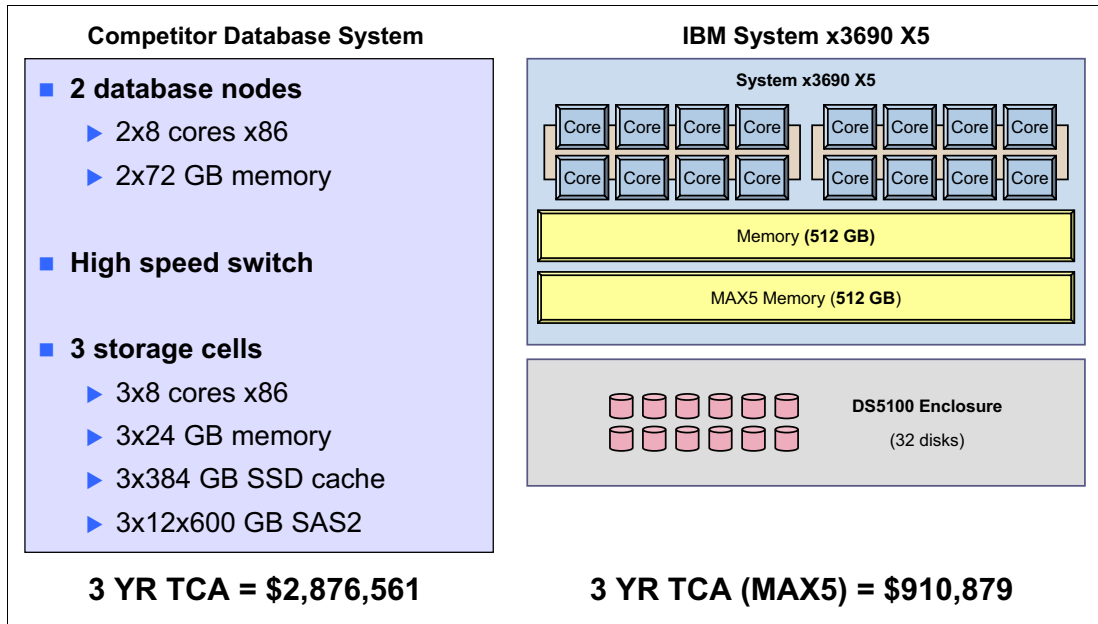


Figure 20 System comparison of competitor versus IBM System x3690 X5

The test in Figure 20 was run with a 1 TB data set, instead of the 1.3 TB data set used in Figure 19 on page 19.

The results in Figure 21 on page 21 are based on IBM internal tests of the System x3690 X5 with a MAX5 solution running a BA workload. This outcome is compared to the results of testing a competitor configuration (previous version; no longer available) with the same workload, both in a controlled laboratory environment. The prices are based on a three-year total cost of acquisition (based on US list prices).

The cost calculation compares the average cost for each report for complex and intermediate report types. The calculation is based on the assumption that the specific types of reports run serially for the three-year life of the machine. The three-year total cost of acquisition includes the expected hardware, software, service, and support. Results might not be typical over a three-year period based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.



Figure 21 DB2 on System x3690 X5 with MAX5 runs serial BA queries faster than a comparable system

Results confirmed that the System x3690 X5 was faster for each report type (complex, intermediate, and simple). The reports were also less expensive. And it became even less expensive for each report over a three-year period.

Integrated solutions: IBM Smart Analytics 5600S R2

Next, we look at a BA System x solution that is an already integrated, modular system optimized for analytics workloads. The solution is a System x server known as the IBM Smart Analytics 5600S R2⁵. This server uses Intel technology. It is optimized for analytics workloads by using IBM warehousing and business intelligence software to transform information into actionable insight. Figure 22 on page 22 describes Smart Analytics 5600S R2.

⁵ For product specifications, see this website:
<http://www-01.ibm.com/software/data/infosphere/smart-analytics-system/5600/>

- A complete, workload optimized, data warehouse system
 - ▶ Including IBM Software, System x hardware, and IBM storage with SSD
 - ▶ System is pre-installed, tuned, and "Data Load Ready"
- Minimize installation, integration, configuration, and tuning requirements

Analytics

- ☑ Cognos 8 Business Intelligence
- ☑ Cubing Services
- ☑ IBM SPSS 19 statistical package
- ☑ Text Analytics and Data Mining

Data Warehouse Software

- ☑ InfoSphere Warehouse
- ☑ Advanced Workload Management
- ☑ Tivoli System Automation

Hardware and Storage

- ☑ IBM System x (System x3500 M3, System x3850 X5, System x3650 M3)
- ☑ System Storage DS3500
- ☑ SSD standard in all configurations



Faster Time To Value

- Deploy in days instead of months
- Focus on creating business value, not installing hardware and software

Costs less than building it yourself

- Less staff and expertise required to implement, tune, and maintain

Unique offering

- Modular for flexibility and growth
- No competitor offers such an integrated analytics solution

Figure 22 Overview of the IBM Smart Analytics system

Additional features of the IBM Smart Analytics 5600S R2 provide the following benefits:

- ▶ Built on a foundation of IBM data warehouse management software, storage, and the System x server platform
- ▶ Supports the Linux operating system
- ▶ Offers a wide range of BA capabilities that include business intelligence reporting, analysis, dashboarding, data mining, cubing services, and text analytics
- ▶ Ensures an affordable upgrade path with a flexible architecture that is designed to address the future growth of data volume, concurrent users, and analytics capabilities
- ▶ Incorporates FusionIO SSD technology in optional modules to improve response time and boost efficiency for complex workloads
- ▶ Accelerates system deployment with the simplicity of an appliance, while retaining the flexibility of a custom integration approach
- ▶ Allows for coordinated updates and maintenance on the control console
- ▶ Includes setup services and a single point of support for convenience

Figure 23 on page 23 depicts the process we used for the next test, conducted in 2011: process for running concurrent operational reports with the IBM Smart Analytics 5600S R2. The workload consisted of 80 users, each having one type of test, all submitting their queries to be run in parallel. As each test was completed, it was submitted again. After four hours, the test ended.

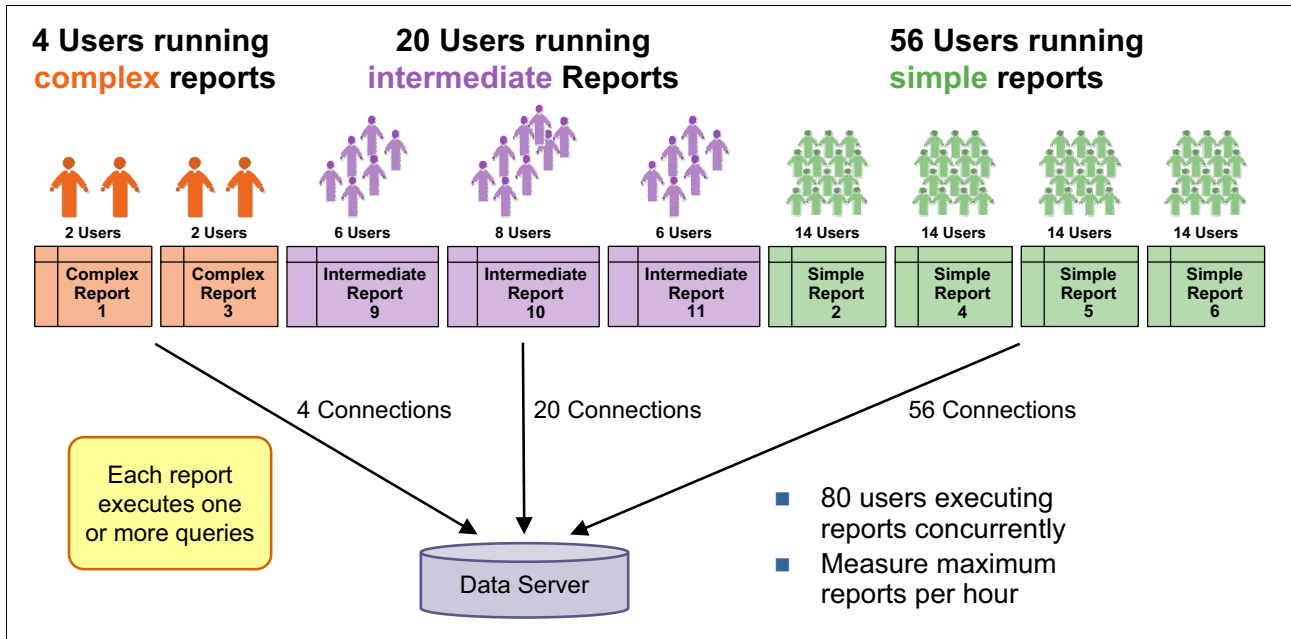


Figure 23 Process for running concurrent operational reports (multiple users)

For each type of report, we calculated the reports generated per hour. Figure 24 on page 24 shows that the Smart Analytics 5600S R2 completed more reports and has a lower three-year total cost of acquisition. The results in Figure 24 are based on IBM internal tests of the Smart Analytics 5600S R2 solution. This outcome is compared to the results of testing a competitor configuration (previous version; no longer available) with the same workload, both in a controlled laboratory environment. The test compares the performance of running different reports of varying complexity, concurrently. The prices are based on a three-year total cost of acquisition (based on US list prices).

Results might not be typical and vary based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.

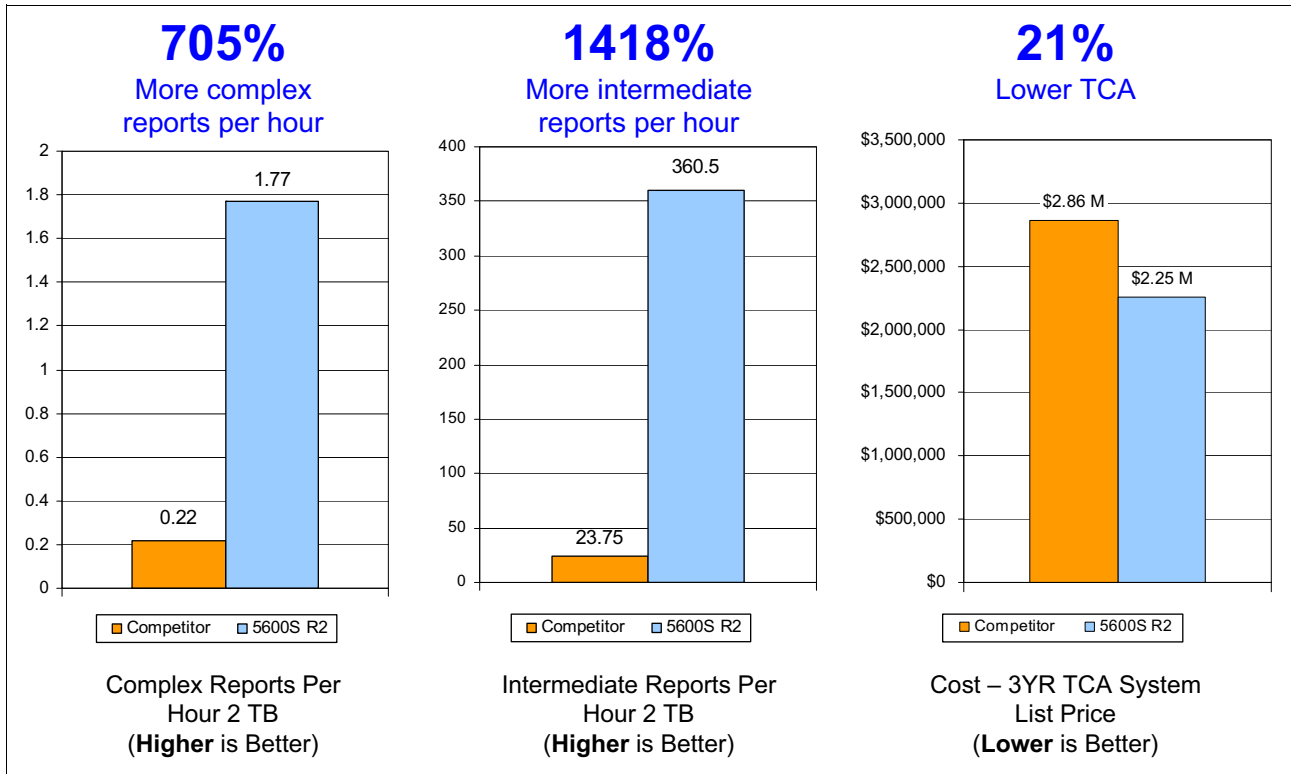


Figure 24 Throughput of complex reports during concurrent workloads

Another performance comparison was conducted that used Smart Analytics 5600S R2. The comparison was conducted as a fixed report study. We ran a specific number of tests, according to specifications listed in Figure 25. These tests were run by 80 parallel users until all were completed.

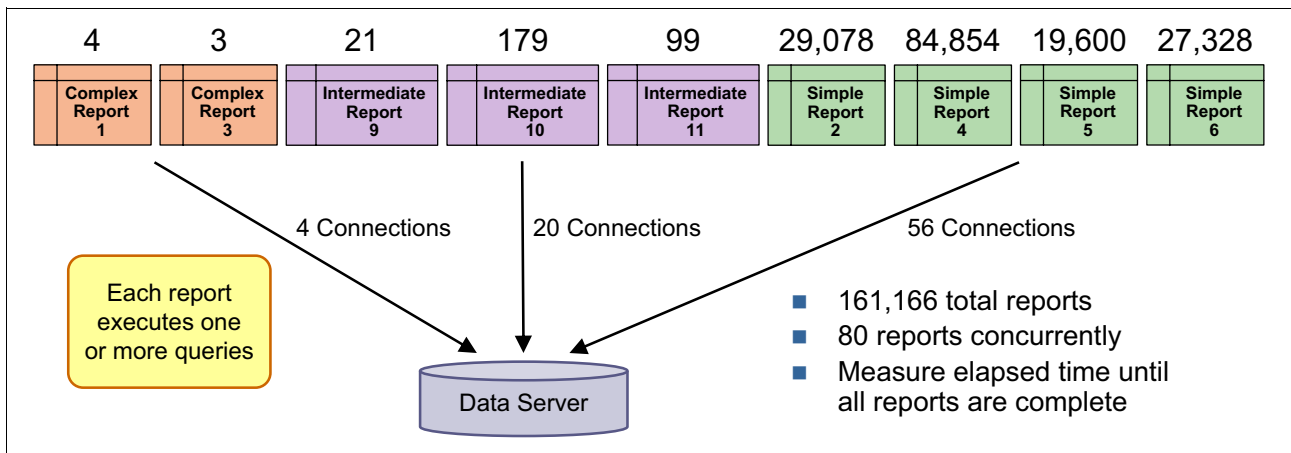


Figure 25 Process for testing parallel execution of reports (fixed execution)

Figure 26 on page 25 shows the results. It took 76% less time for the Smart Analytics 5600s R2 to complete the reports than it did for the comparison server. Running the reports was accomplished at a 21% lower cost over a three-year total cost of acquisition. Values in Figure 26 are based on IBM internal tests of the Smart Analytics 5600S R2. The Smart Analytics solution is compared to the testing of a competitor configuration (previous version; no longer available) with the same workload, both in a controlled laboratory environment.

The three-year total cost of acquisition includes hardware, software, service, and support. The values are based on US list prices. The test compares the performance of running different reports of varying complexity, concurrently. Results might not be typical and vary based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.

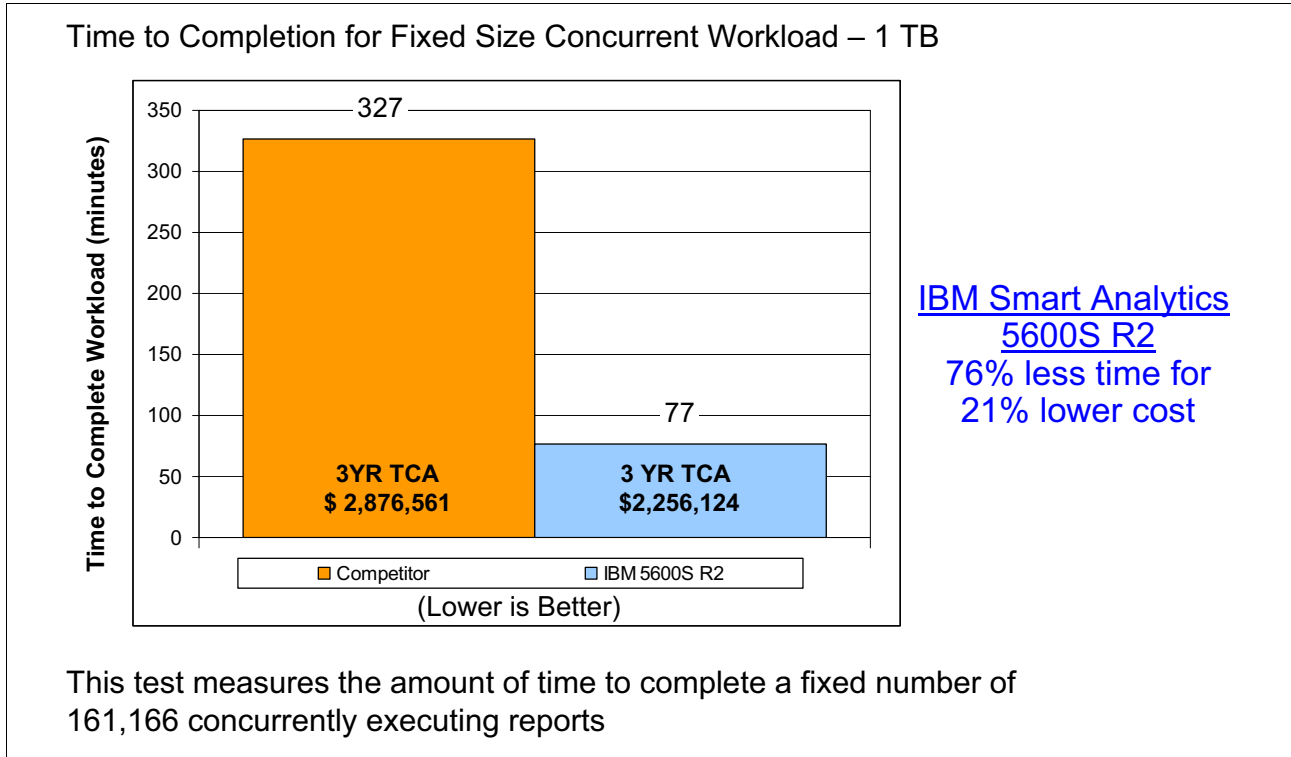


Figure 26 In this study, we assessed the time to run a fixed set of concurrent reports

Appliance solutions: Netezza

Figure 27 on page 26 provides a brief tour of the Netezza 1000 to describe what makes it so unique.

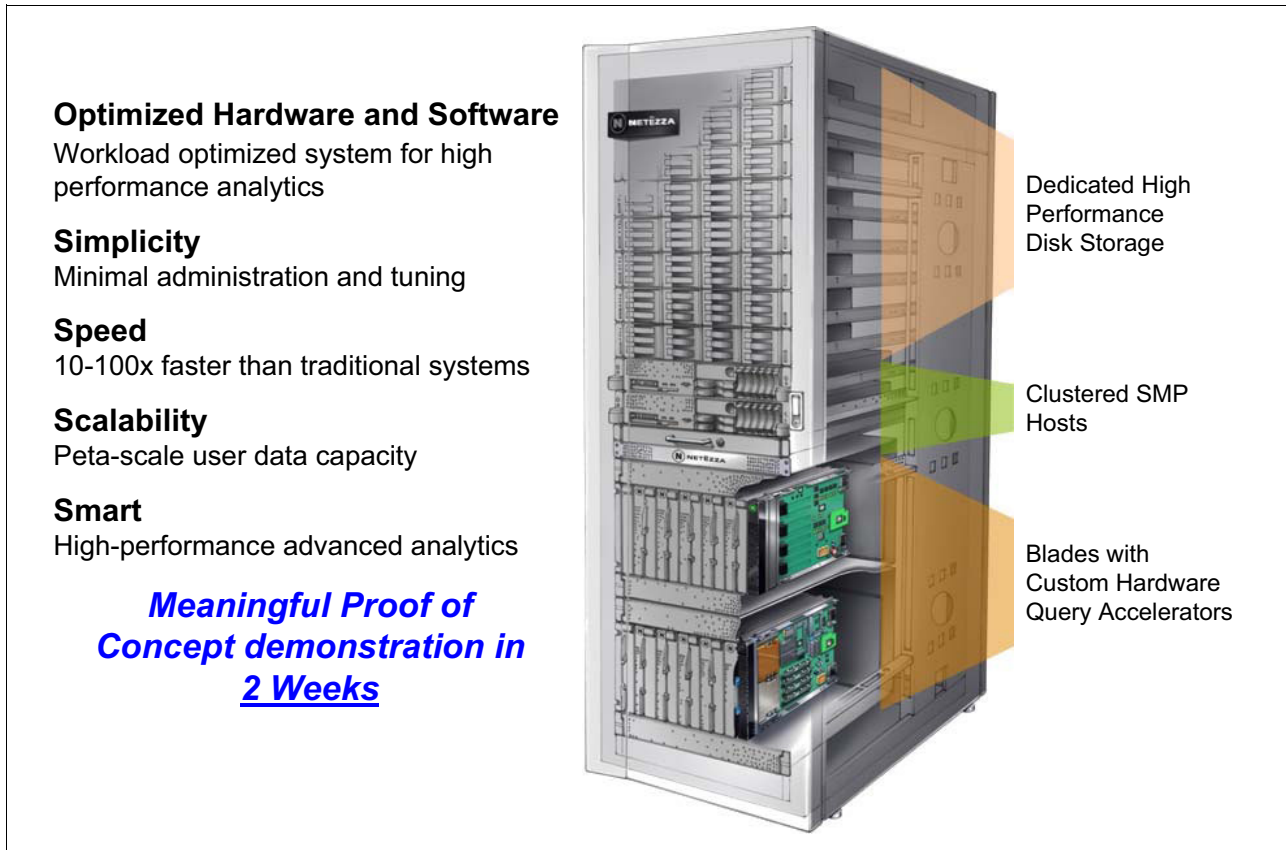


Figure 27 Overview of the IBM Netezza 1000 data warehouse appliance

The following additional benefits are unique to the IBM Netezza 1000:

- ▶ The hardware and software are optimized and fine-tuned to process analytic queries; therefore, the appliance requires zero to minimal tuning for performance.
- ▶ The software is designed to push the absolute majority of processing to the Massively Parallel Processing (MPP) network.
- ▶ The architecture maximizes processing efficiency by ensuring there are no bottlenecks along the I/O path.
- ▶ The unique feature of the Netezza 1000 is the hardware-based query acceleration, which filters out 95 - 98% of the data that is not needed for query processing. The benefit is that the query acceleration manages the deluge of data typical of data warehousing deployments.

The Netezza 1000 is suitable for serial queries for the following reasons:

- ▶ Purpose-built architecture for BA:
 - Shared nothing* architecture for MPP
- ▶ Data streaming and hardware-based acceleration, based on Field Programmable Gate Array (FPGA), speeds up queries in the following ways:
 - Data streaming moves data from a disk by using an FPGA chip to the processors without network latency and unnecessary processor usage.
 - Each MPP node is self-contained and has minimal data movement.

- ▶ Faster specialized code path:
 - Purpose-built optimizer is tailored for complex analytics with higher predictability.
 - The absence of traditional database structures and OLTP features results in a simpler code path and faster performance.

For the next test, to measure performance levels, we compared the Netezza 1000 with a comparison machine. Results and product specifications are noted in Figure 28, where we compared these two machines against the 1 TB data set, serial execution test.

Values in Figure 28 are the list prices that were published on the IBM website (<http://www.ibm.com>) when the analysis was conducted in mid-2011. Values are based on a three-year total cost of acquisition of a competitor configuration (previous version; no longer available). Total cost of acquisition includes the hardware, software, service, and support that are used to run analytic workloads. Prices are based on US list prices and vary by country. Users of this document need to verify the applicable data for their specific environment.

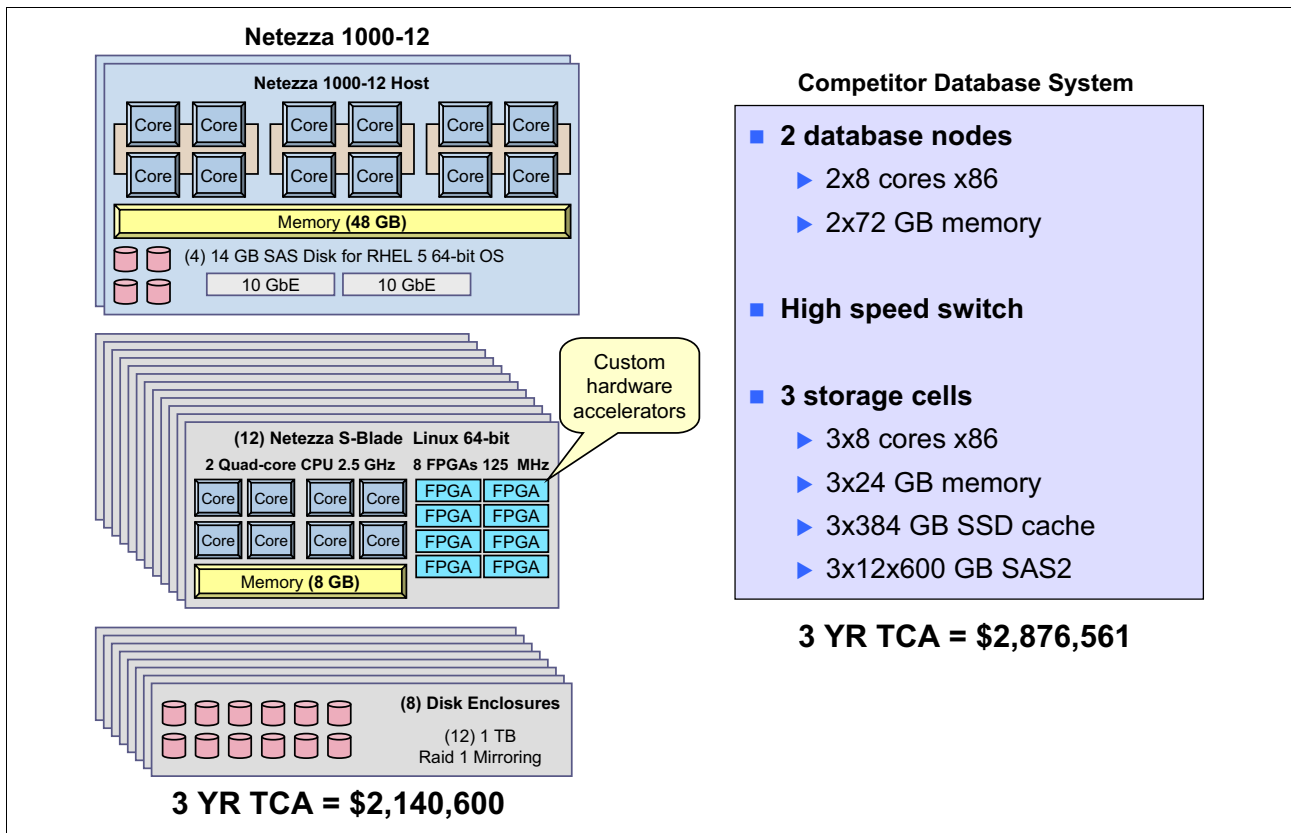


Figure 28 System comparison of IBM Netezza 1000-12 versus competitor

Additional product specifications of the IBM Netezza 1000 include the following benefits:

- ▶ High-performance analytics with its default settings: Requires no indexing or tuning, leading to shorter deployment cycles and faster time-to-value
- ▶ A powerful platform, designed specifically for unifying business intelligence and advanced analytics on large volumes of data
- ▶ Appliance simplicity: Easy to deploy and manage, and dramatically simplifies your data warehouse and analytic infrastructure
- ▶ Powerful scaling and low total cost of ownership: Scales from 1 TB to 1.5 petabytes, delivering 10 - 100 times the performance improvements over traditional architectures

- Unique asymmetric massively parallel processing (AMPP) architecture that combines open, IBM blade servers, and disk storage with the IBM patented data filtering FPGAs

Figure 29 shows the results of the response times for running the analytic queries. Not only was the Netezza considerably less expensive (as shown in Figure 28 on page 27), it was also much faster than the comparison machine. Values in Figure 29 are based on IBM internal tests of the Netezza 1000-12 solution. The Netezza solution is compared to the testing of a competitor configuration (previous version; no longer available), running analytics workloads in a controlled laboratory environment.

Three-year total cost of acquisition includes hardware, software, service, and support. Values are based on US list prices. The test compares the performance of running different reports of varying complexity. Results might not be typical and vary based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.

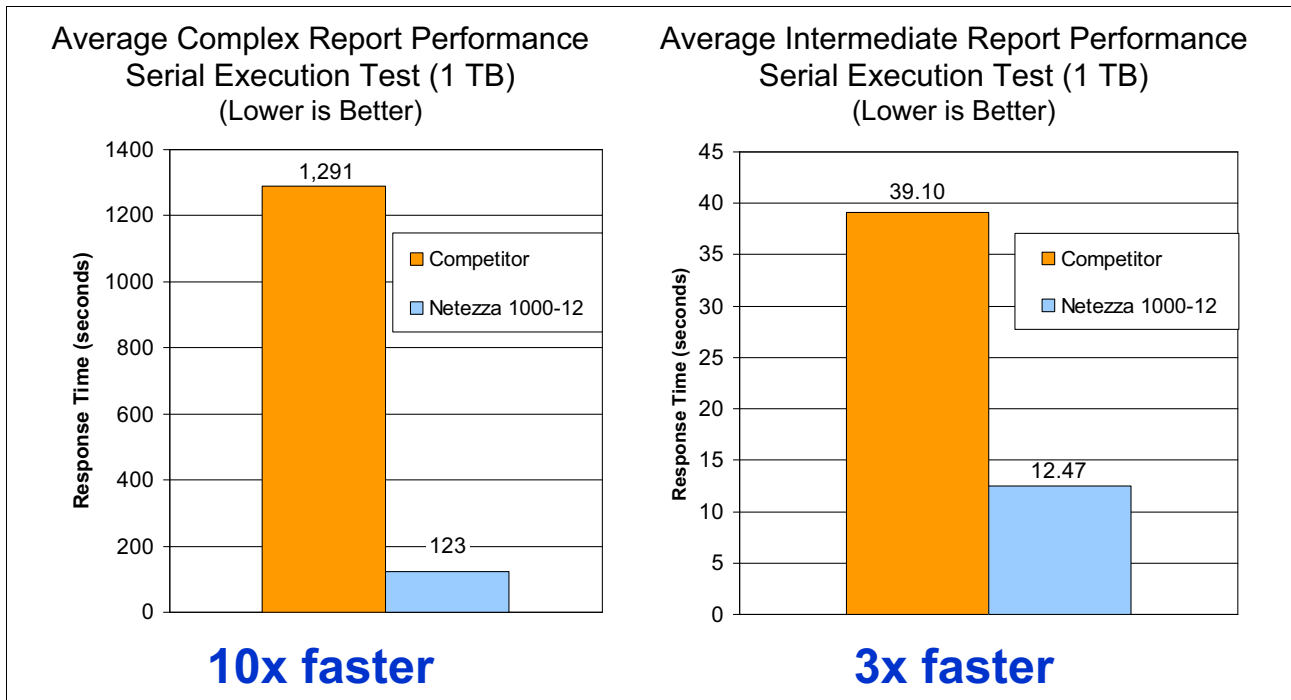


Figure 29 Response times for running analytic queries

Figure 30 on page 29 shows test results that compared the Netezza 1000 and a comparison machine to determine the performance of the cost per report over a three-year period. As shown in Figure 30, results showed that Netezza costs considerably less per report. Values in Figure 30 are based on IBM internal tests of the Netezza 1000-12 solution, conducted in 2011. The Netezza solution is compared to the testing of a competitor configuration (previous version; no longer available), running analytics workloads in a controlled laboratory environment.

Three-year total cost of acquisition includes hardware, software, service, and support. Values are based on US list prices. The test compares the performance of running different reports of varying complexity, concurrently. Results might not be typical and vary based on actual configuration, applications, specific queries, and other variables in a production environment. Users of this document need to verify the applicable data for their specific environment.

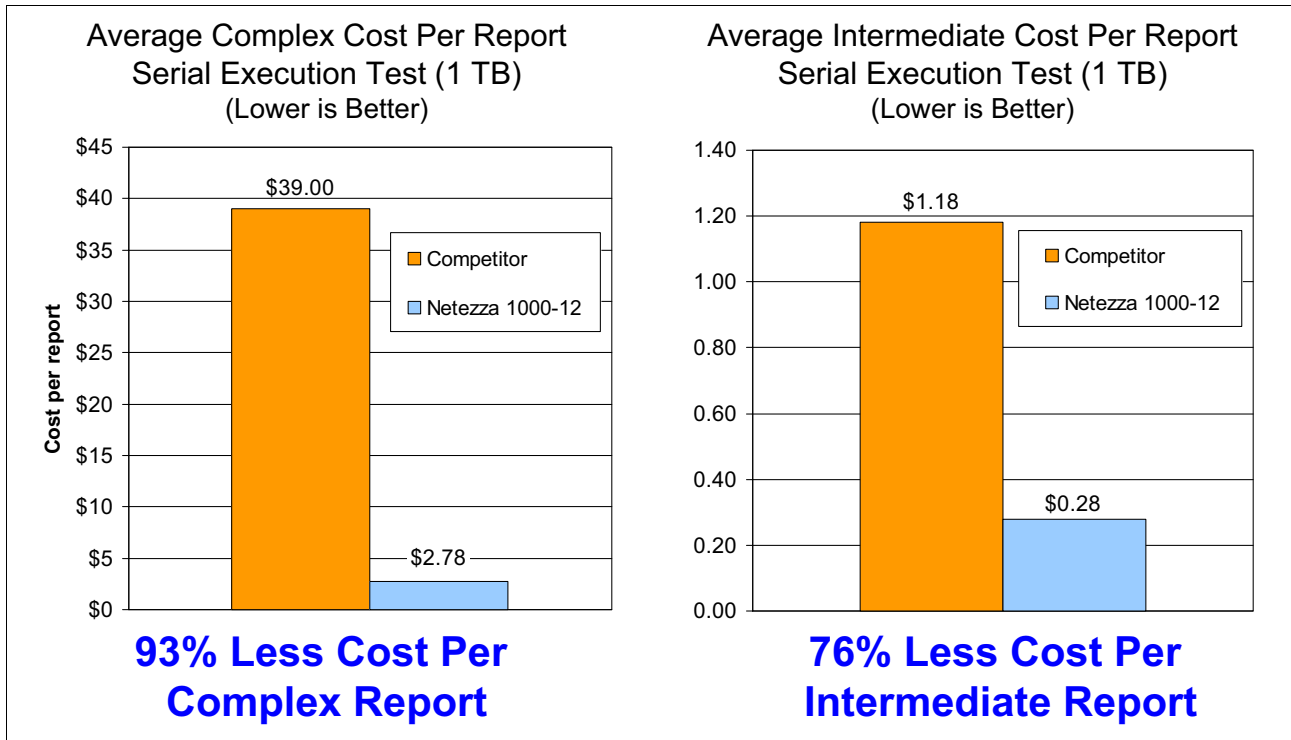


Figure 30 Comparison of cost per report

The entire study, shown in Figure 30, was completed in a span of four hours for Netezza. However, it took about four days of reading manuals and tuning (ASM rebalance, partitioning, and parallelizing hints) for the comparison database to obtain optimal performance.

Figure 31 on page 30 describes the performance of Netezza 1000 (with no tuning) and a comparison machine (after four days of tuning).

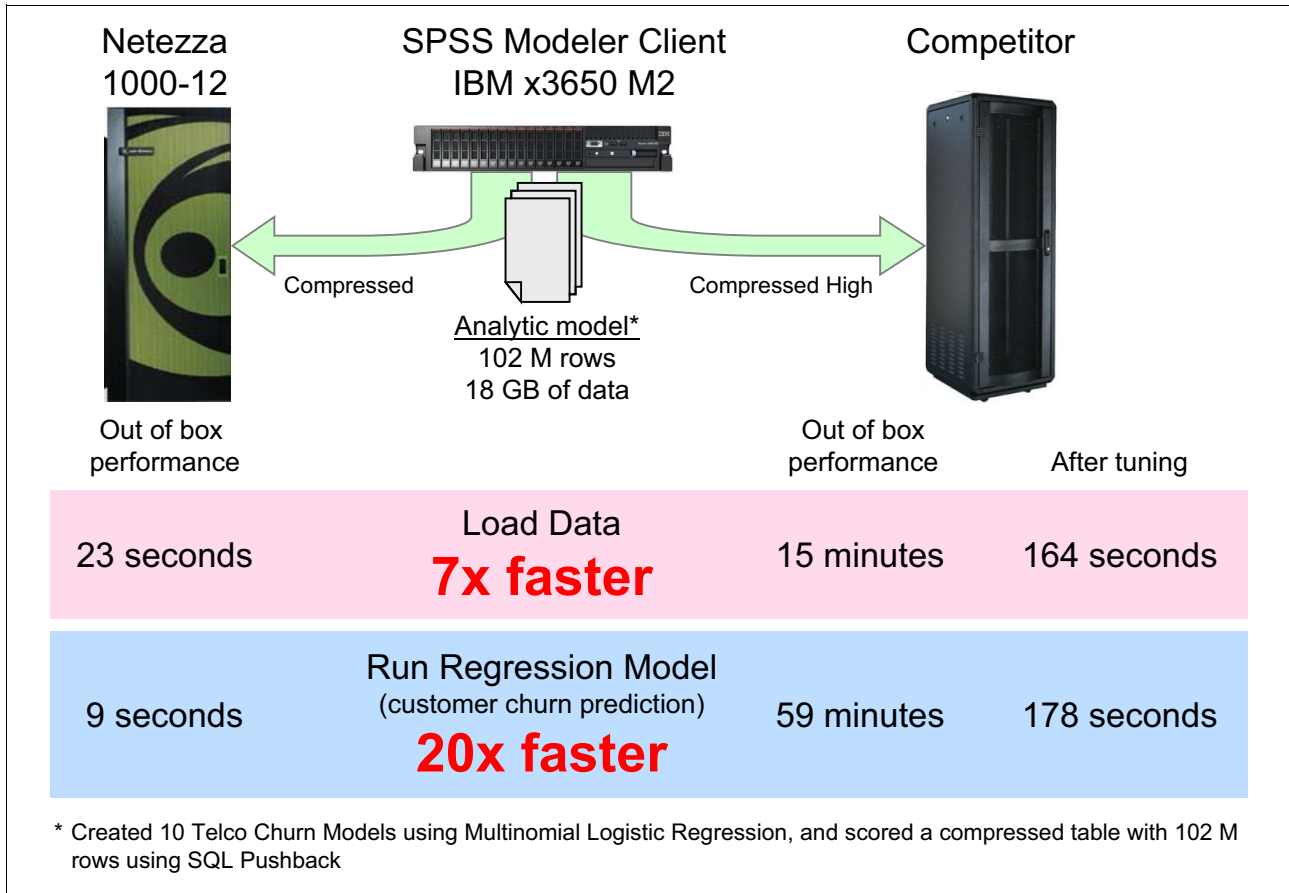


Figure 31 Right from the beginning, Netezza 1000 requires no tuning and outperforms a comparable system

A description of the test method, shown in Figure 31, follows.

Data Load Phase

The data was from a Telco study that used a table with 102 million rows.

Compression was used for both databases (Netezza supports only compressed data, and we used Compress Query Low and High for the comparison database). The comparison database Compress Query High data load is slower than Compress Query Low. We used the better number to provide the comparable system the benefit of the doubt (for scoring performance, the number does not change between Compress Low or Compress High).

Performance of the built-in components indicates that the tests were conducted with the same SQL queries (that is, there was no tuning on either system). Netezza did not require tuning, and the data definition language (DDL) included random partitioning. For the comparison database, performance of the default settings was conducted for 15 minutes (Compress Query High) and five minutes (uncompressed). We tuned the comparison database manually, which created four table spaces (each for input file S100M and output file OP100M). And we used 64 partitions and 64 parallel threads (degree) and modified INSERT SQL with parallelization hints (such as /*+ APPEND */).

Table 3 on page 31 shows the data load times with compression time (seconds) for this test.

Table 3 Data load times and compression time (seconds)

| | Netezza | Comparison database (Normal compression) | Comparison database (High compression) |
|---------------------------------|--------------|--|--|
| Create table S100M | 0.246 | 0.48 | 0.33 |
| Load a 100,000-record data file | 3.315 | 10.47 | 18.01 |
| Fold (ten times) | 19.809 | 102.63 | 145.55 |
| Total time | 23.37 | 113.58 | 163.89 |

Table 3 shows the results of this test in terms of total time, where Netezza performed eleven times faster than the comparison machine. The Netezza was tested with its default settings. To be ready for the test, the comparison machine was configured over a period of four days.

Figure 32 describes how comparison systems are sometimes designed to support the “one-size-fits-all” concept. This Redpaper publication shows that this approach is typically not the best solution. Our workload-optimized solutions can offer more efficiency and better price performance.

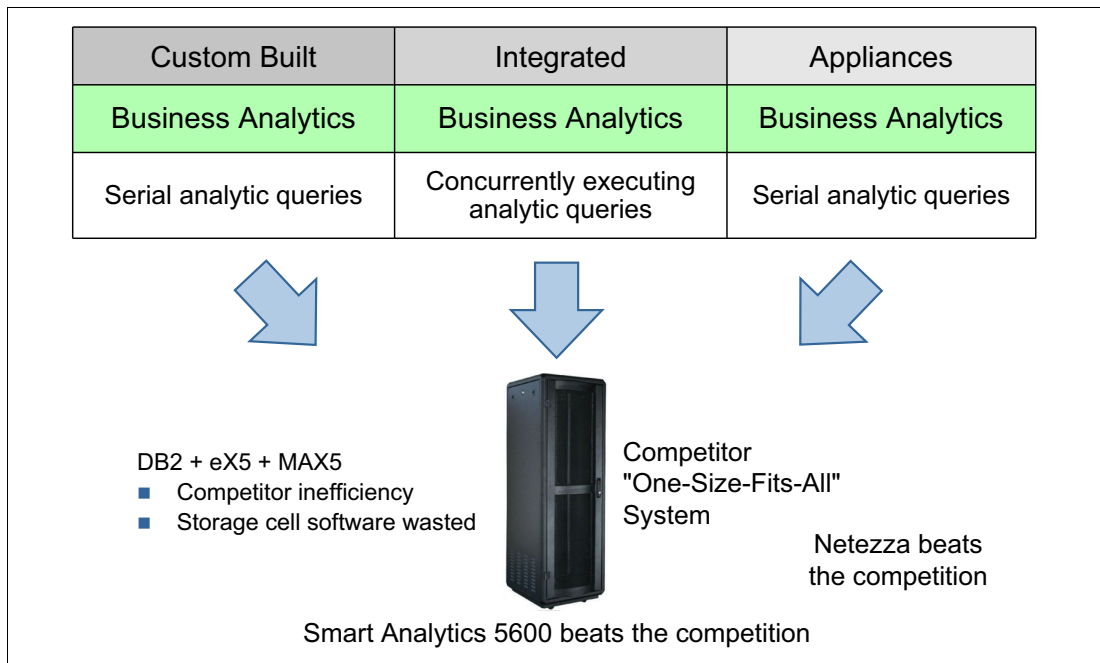


Figure 32 Database workloads can benefit from one of the IBM workload-optimized solutions

IBM software and System x are best on Intel for data workloads

For OLTP and BA workloads, this paper describes how IBM software and System x can meet your needs, without regard to complexity. One size does not fit all; IBM is aware of that fact. Part of the success of the IBM software and System x is the powerful, expansive Intel processor. Together, IBM software, System x, and Intel are designed to address your database workload requirements. Figure 33 on page 32 defines the primary factors that make IBM and Intel a good combination for your OLTP and BA workloads.

- OLTP workloads and the competition:
 - ▶ DB2 and System x with eXFlash provided
 - 54% more throughput at up to 46% less cost
 - ▶ DB2 and System x with Storwize V7000 Easy Tier provided
 - 1.25x more throughput at 60% less cost
- Business Analytics workloads and the competition:
 - ▶ DB2 and System x with MAX5 won for serial reports in performance and cost
 - ▶ IBM Smart Analytics System 5600 R2 won for concurrent workloads in performance and cost
 - ▶ Netezza won for serial reports in performance and cost with an appliance solution

Figure 33 IBM software and System x are best on Intel for OLTP and BA workloads

Summary

Two specific database workloads are discussed in this Redpaper document: OLTP and BA workloads.

For OLTP workloads, the eX5 family of servers offers the high performance processors for CPU processing. When I/O throughput becomes a bottleneck, System x offers eXFlash internal SSDs to boost I/O. When more boost is needed, the Storwize V7000 with Easy Tier provides it.

For BA workloads, choose from the most configurable high-end systems with the best memory, or choose pre-configured systems (with several options) that are tuned for BA. You can also select an appliance that does what you need and works quickly, with its initial default settings.

For both workload types, IBM offers the hardware and software to optimize your database workloads.

For more information

For more information about the products described in this publication, see these websites:

- ▶ IBM System x overview
<http://ibm.com/systems/x/>
- ▶ IBM eX5 enterprise systems overview
<http://www.ibm.com/systems/info/x86servers/ex5/>
- ▶ IBM System x3690 X5 server overview
<http://www.ibm.com/systems/x/hardware/enterprise/x3690x5/>
- ▶ *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and IBM BladeCenter HX5*, REDP-4650, December 2011.
<http://www.redbooks.ibm.com/abstracts/redp4650.html>

- ▶ IBM System x3850 X5 overview
<http://www.ibm.com/systems/x/hardware/enterprise/x3850x5/>
- ▶ IBM BladeCenter HX5 overview
<http://www.ibm.com/systems/info/x86servers/ex5/blade/>
- ▶ IBM Easy Tier overview
http://www.ibm.com/systems/storage/disk/storwize_v7000/
- ▶ IBM eXFlash storage overview
<http://www.ibm.com/systems/info/x86servers/optimized/database/index.html>
- ▶ IBM MAX5:
 - *Add Memory, Improve Performance, and Lower Costs with IBM MAX5 Technology*, 2012, REDP-4846.
<http://www.redbooks.ibm.com/abstracts/redp4846.html>
 - *The MAX5 Advantage: How IBM System x MAX5 Benefits Microsoft SQL Server Data Warehouse Workloads*, REDP-4751, 2011.
<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 ProtecTIER® overview
<http://www.ibm.com/systems/storage/tape/ts7650g/>
- ▶ IBM Systems overview
<http://www.ibm.com/systems/>

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. Prior to joining STL in 2010, Tim was the Global Community of 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 Bachelor and Master degrees in Computer Science from USC. He is a yearly guest lecturer at USC and UCLA.



Karen Lawrence is a Technical Writer with the IBM ITSO team in North Carolina, US. She has 25 years 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. Karen holds a Bachelor 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, IBM Product Marketing - High End System x, for his contributions and sponsorship of this project.

Thanks also to Linda Robinson, IBM Graphics Specialist; and to the IBM ITSO editing team 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 client 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.

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

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-4848-00 was created or updated on August 31, 2012.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
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 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:

BladeCenter®
Cognos®
DataPower®
DB2®
Easy Tier®

IBM®
ProtectTIER®
Rational®
Redbooks®
Redpaper™

Redbooks (logo) ®
Storwize®
System x®
Tivoli®

The following terms are trademarks of other companies:

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

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

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