IBM eserver

Understanding IBM eserver xSeries Benchmarks

What industry benchmarks are used in the xSeries marketplace

How to interpret the results of the benchmarks

How the benchmarks relate to client configurations

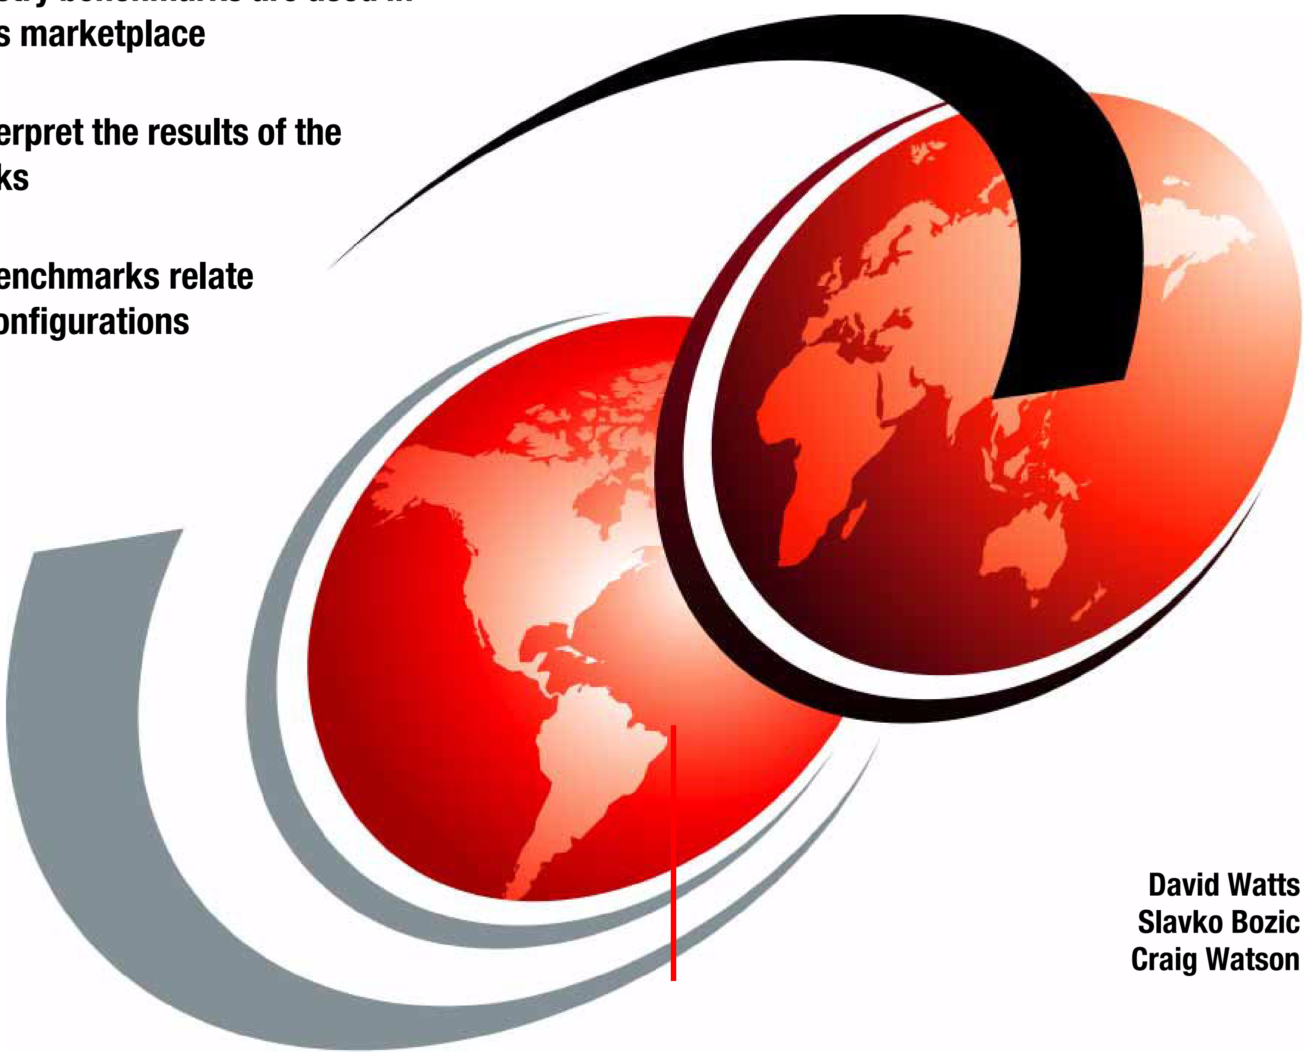

David Watts
Slavko Bozic
Craig Watson

Redpaper

ibm.com/redbooks

International Technical Support Organization

Understanding IBM eserver xSeries Benchmarks

May 2005

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (May 2005)

This edition applies IBM @server xSeries and BladeCenter servers.

This document created or updated on May 13, 2005.

Contents

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.

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

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 illustrates 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. 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 IBM's application programming interfaces.

Trademarks

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

eserver®	AIX 5L™	Perform™
eserver®	AIX®	PowerPC®
ibm.com®	BladeCenter™	POWER™
iSeries™	Domino®	Redbooks (logo) ™
i5/OS™	DB2®	ServerProven®
pSeries®	IBM®	Tivoli®
xSeries®	Lotus®	TotalStorage®
z/OS®	Notes®	VSE/ESA™
z/VM®	OpenPower™	WebSphere®
zSeries®	OS/400®	X-Architecture™

The following terms are trademarks of other companies:

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

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

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

TPC benchmark and TPC-C are certification marks of the Transaction Processing Performance Council.

SPEC benchmark is a certification mark of the Standard Performance Evaluation Corporation.

Preface

Many models of the IBM @server® xSeries® family maintained a leadership position for benchmark results for several years. These benchmarks help clients position xSeries servers in the marketplace, but they also offer other advantages to clients including driving the industry forward as a whole by improving the performance of applications, drivers, operating systems, and firmware.

There is a common misconception that industry benchmark results are irrelevant because they do not reflect the reality of client configurations and the performance and transaction throughput that is actually possible in the "real world". This Redpaper shows that benchmarks are useful and relevant to clients and that benchmark results are useful when attempting to understand how one solution offering performs over another.

The purpose of this Redpaper is to explain what benchmarks are and how to interpret the benchmark results so you can understand how they relate to their own server plans. The major industry benchmarks from the Transaction Processing Performance Council (TPC) and the Standard Performance Evaluaction Corporation (SPEC) are described, explaining how they relate to specific client application types.

This paper is for clients, IBM® Business Partners, and IBM employees who want to understand benchmarks on xSeries servers.

The team that wrote this Redpaper

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

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks on hardware and software topics related to IBM @server xSeries systems and associated client platforms. He has authored over 30 redbooks and redpapers, most recently the IBM Redbook, *Tuning IBM @server xSeries Servers for Performance*. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 15 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.

Slavko Bozic is an xSeries Technical Specialist and Technical Advisor in Sweden. He has eight years of experience in the IT industry. He holds a degree from the Programme in Automatic Data Processing at the University of Gothenbourg (Department of Informatics). His areas of expertise include server and storage consolidation, Citrix and VMware. He has been working for Pulsen, an IBM Business Partner in Sweden for eight years.

Craig Watson is an xSeries Technical Specialist in New Zealand. He has nine years of experience in the IT industry, and has extensive experience in Windows® and UNIX® performance tuning. He holds a Masters degree in Electrical Engineering from the University of Auckland and has worked at IBM for three years.

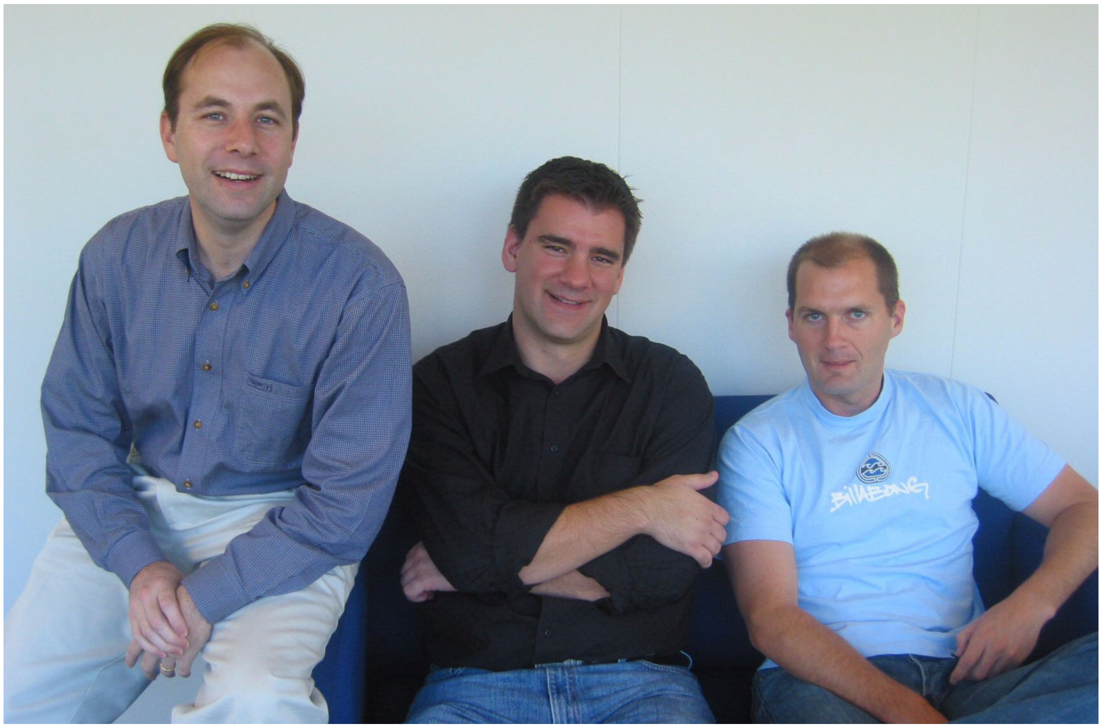

The team (l-r): David, Slavko, Craig

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

Tamikia Barrow
Margaret Ticknor
Jeanne Tucker
International Technical Support Organization, Raleigh Center

Chitra Balachandran
Matt Eckl
Chris Floyd
Joe Jakubowski
Tricia Hogan
Phil Horwitz
Doug Pase
Tricia Thomas
IBM xSeries Performance Lab, Raleigh

Jay Bretzmann
IBM xSeries Marketing Management

Become a published author

Join us for a two-to-six week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will team with IBM technical professionals, Business Partners, or clients.

Your efforts help increase product acceptance and client satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HZ8 Building 662
  P.O. Box 12195
  Research Triangle Park, NC 27709-2195

1

Benchmarks 101

A benchmark is a standardized problem or test used to measure system performance. The purpose is typically to make some sort of comparison between two offerings, whether they are software, hardware, or both.

There are many types of benchmarks undertaken, varying dramatically in purpose, size, and scope. A very simple benchmark may consist of a single program executed on a workstation that tests a specific component such as the CPU. On a larger scale, a system-wide benchmark may simulate a complete computing environment, designed to test the complex interaction of multiple servers, applications, and users. In each case the ultimate goal is to quantify system performance to take measurements.

This chapter contains the following topics:

1.1 Types of benchmarks

There are three overall types of benchmarks:

- Industry-standard benchmarks

  This Redpaper concentrates on industry-standard benchmarks.These are many well known benchmarks developed, maintained, and regulated by independent organizations. The benchmarks are designed to represent client workloads (for example, e-commerce or OLTP) and allow the reader to make comparisons between systems when the workload matches their intended use. The configurations are based on off-the-shelf hardware and applications. We introduce industry benchmarks in 1.4, "Current industry standard benchmarks" on page 4.

- Unregulated benchmarks

  Also very common are unregulated benchmarks that are application or component specific. You can sometimes purchase or download these benchmark suites to test how specific components perform.

  Use extreme caution when using these tools. By far the majority of testing suites available test workstation performance and are not relevant for testing server performance. See "Server versus workstation benchmarks" on page 5 for more information.

- Client workload benchmarks

  A third category of benchmarking involves benchmarking with a client's actual workload. This yields the most relevant information, but is difficult to do. We discuss this type in detail in Chapter 5, "Client-related benchmarks" on page 49.

> **Important:** In order to test the performance of a server, or to use published benchmark results to compare systems, you must understand the characteristics of the intended workload. Server performance differs with different workloads. A server that produces industry-leading performance under one workload, may perform poorly under another.
>
> Performance figures from a benchmark or test that do not reasonably resemble the intended workload of a server has limited meaning.

The only accurate method of determining how a server will perform under a particular workload is to test with that environment. Unfortunately this can be very difficult and expensive for clients to do. For this reason, vendors such as IBM spend a considerable amount of time and money performing industry benchmarks on their servers with different applications and workloads. From this, performance information is produced that clients can use to help make informed decisions.

The following sections describes why industry standard benchmarks are necessary, how you can use the benchmarks, and limitations of which to be aware.

Following are some of the many factors that influence performance:

- The basic configuration model
- The database
- The disk subsystem
- The memory subsystem
- Especially the way the application is configured and utilized

The results of any benchmark, no matter how well it simulates a real-world scenario, cannot be confidently translated to a user's environment.

The published results of industry standard benchmarks are extremely useful *approximates* rather than *definitive*, guides to relative performance, if real workload testing is not practical.

1.2 A level playing field

How can you compare and evaluate the performance claims of one vendor against those of another? In the early 1980s it was generally recognized that metrics such as millions of instructions per second (MIPS) were inadequate for gauging the performance of a server executing an Online Transaction Processing (OLTP)-type workload. Hardware vendors at the time began quoting the performance of their systems in terms of transactions per second (tps), but they revealed few details of the tests used to produce such figures. This made it extremely difficult to make informed comparisons between different systems.

In 1984, a new benchmark was proposed that was intended to be vendor neutral. The requirements of the benchmark were described at a high functional level, rather than distributed as an executable program. This was an important distinction in contrast to other benchmarks at that time, because it allowed the test to be implemented on any type of hardware or software. The development of this benchmark was essentially the beginnings of the Transaction Processing Performance Council (TPC), an independent organization that is responsible for creating effective benchmarks, and for ensuring that the rules of the benchmarks are followed.

At around the same time, a small group of workstation vendors established another non-profit organization to create and monitor standardized CPU benchmarks. Both this organization, named the Standard Performance Evaluation Corporation (SPEC), and TPC developed into the most recognized and widely accepted industry standardization bodies for performance tests.

You can find information regarding these organizations the following Web sites:

http://www.tpc.org/
http://www.spec.org/

1.3 A rigorous process means credible results

Industry standard benchmarks are performed in accordance with a tightly defined set of rules. These rules are agreed upon amongst participants, and a mechanism for ensuring that the rules are followed is put in place. In the case of the TPC benchmarks, an independent auditor reviews and verifies the benchmark result before it can be released.

A full disclosure report is also published that details all components of the benchmark, including the hardware, software, and every parameter setting required to reproduce the result. This process enables competitors to examine how the performance figure was achieved, learn tuning techniques, and challenge the result if any rules were broken.

The degree of scrutiny to which these benchmark results are subjected means that they are extremely credible. Although the actual performance figures reached are generally not realistic in a production environment, the results demonstrate the relative strengths of a system architecture under that type of workload.

A further discussion about the industry benchmark process can be found in 3.2, “Rules of engagement” on page 21.

1.4 Current industry standard benchmarks

There are a large number of industry standard benchmarks available, the TPC and SPEC maintain the most widely recognized and accepted ones. Benchmarks that are specific to an application are also often quoted, for example Oracle Applications Standard Benchmark and SAP Standard Application Benchmark. Table 1-1 shows a number of the current industry standard benchmarks available, and the nature of workloads that they represent.

These benchmarks are discussed further in Chapter 3, “Industry benchmarks” on page 19.

Table 1-1 Common benchmarks and the workloads they represent

Benchmark	Type of workload
TPC-C and TPC-E	Online transaction processing
TPC-H	Ad-hoc decision support
TPC-W and TPC-App	Transactional Web e-Commerce
SPECweb	Web serving static and dynamic pages
SPECweb_SSL	Web serving with Secure Sockets Layer (SSL)
SPECjbb	Server-side Java™
SPECjAppServer	J2EE-based application server
SPEC HPC	High performance computing, CPU, interconnect, compiler and I/O subsystems
SPEC CPU	Compute intensive, integer and floating point performance
Oracle Applications	Models the most common transactions on the seven most used Oracle Application modules
SAP Standard Application	Suite of benchmarks for mySAP Business Suite.
BaanERP	Transaction processing environment of iBaan ERP applications
Notesbench	Simulates Domino® workstation-to-server or server-to-server operations
Exchange MAPI Messaging	Measures the maximum messaging throughput of a Microsoft® Exchange Server
Linpack HPL	Solving a dense system of linear equations Used to compile the top 500 supercomputer list.

1.4.1 Evolving benchmarks

Industry standard benchmarks are generally designed to simulate realistic application workloads. However, as business and technology evolves, application workload characteristics change as a result. A benchmark that successfully replicated a production workload five years ago, may have little resemblance to the workloads that clients use today. For this reason, new benchmarks are constantly being developed that more accurately represent production environments.

An example of this is TPC-C, regarded by many as the industry’s premier benchmark. Although this benchmark effectively highlights strengths in system design, most accept that it no longer represents configurations that clients purchase. Using the current generation of processors, a 4-way server requires hundreds of disks to achieve the throughput required for the central processing units (CPUs) to reach their maximum capacity. This makes the benchmark very expensive to run, and with further advances in processor technology the situation becomes even more difficult.

For this reason, TPC-E is currently in development and is likely to become an alternative to TPC-C. Developers of the TPC-E benchmark aim to significantly reduce the amount of disk required, and make the benchmark easier to run by supplying parts of the code. There is a more realistic schema and Redundant Array of Independent Disks (RAID) protection is encouraged for the data and log, making the benchmark more representative of real environments.

1.5 Invalid comparisons

This section briefly describes instances of which to be aware, where benchmark results are not a valid means to compare performance.

1.5.1 Server versus workstation benchmarks

Use a great deal of caution when you attempt to benchmark servers. Unfortunately it is a difficult and complex task to develop the software and test beds required to reproduce a production-like, multi-user application workload. When testing a server, a mistake frequently made is to copy a large file, or run a simple tool downloaded from the Internet. These are not adequate methods for measuring server performance and can be completely misleading.

Servers and workstations are designed for very different purposes. A workstation performs a single task as quickly as possible. In contrast, servers are generally optimized to service multiple users performing many tasks simultaneously. Many benchmarking tools are designed to test workstations and to stress system components by executing a single task or a series of tasks. A server often performs poorly in such tests as this is not what they are optimized to do.

As an example let us use a benchmarking tool that tests memory bandwidth. This is a common test for workstations. High memory bandwidth is only sustained by an application reading sequentially through memory. For this task, the chipset and processor are programmed to prefetch multiple cache lines on each request so that sustained processor to memory bandwidth can approach about 75% of the bus speed. However, doing this greatly reduces random memory latency since extra cache lines are prefetched and often discarded.

1.5.2 Apples-to-apples comparison

When making a comparison between two different systems, it is extremely important to make an apples-to-apples comparison, meaning that the systems compared are configured in the same way. For example, both systems use the same speed processors with the same amount of cache. It is not uncommon for vendors to make performance claims using results that do not make a fair comparison. Chapter 4, “Understanding benchmark results” on page 39 covers this topic in greater detail.

1.5.3 Performance versus price/performance

The results of industry benchmarks are often stated in one of two different ways, a raw performance figure or a price/performance figure. The rules of many industry standard benchmarks require that the total cost of the system is published as well as the ultimate performance result (transactions per second, for example), so it is a simple matter to divide the price of the system with the performance result achieved to obtain this price/performance number.

One vendor may choose to maximize the performance number at the expense of the price of the system. In this situation, the most and best hardware and software is used to obtain the fastest possible result, often at considerable expense. In contrast, another vendor may choose to minimize the price per performance result by carefully balancing the cost of the hardware and of the performance result achieved.

It is not valid to compare a performance-maximizing benchmark result with a price/performance-leading result.

1.5.4 Results versus reality

The IBM xSeries performance lab produced many industry leading benchmark results on xSeries hardware. These results demonstrate the capabilities of the server and validate the system designs. They are also a valuable source of information to use in a decision-making process. It is important to understand, however, that the actual performance figures produced from these benchmarks are unlikely to be obtained in a real production environment.

The reasons for this disparity include the following:

- System utilization often runs at levels that are not reasonable in a production environment. In most commercial environments, a system operating at a utilization level of very close to 100% cannot handle peaks in demand or allow for future growth.
- The hardware configurations may not be representative of a real production server. For the TPC-C benchmark, hundreds of disks feed as much data as possible to memory and the CPUs.
- In addition to a large number of drives, each drive is configured with a “stroke” (portion of the disk’s surface actually containing data) of less than 15%. The objective is to minimize latency by reducing the disk drive head movement as much as possible. This means only a small portion of each disk is used. This is unrealistic in a production environment.
- Many benchmarks do not require RAID arrays to protect data. For these benchmarks data is striped across the disks in an unprotected manner (that is, RAID-0) to maximize performance and minimize cost.
- Servers respond differently to different workloads. A server that is industry leading in one benchmark may perform poorly in another. When using published benchmarks as a guide to selecting a server, the results of benchmarks that most closely resemble the intended workload should be examined.
- A vendor benchmark team typically has more time and greater access to highly skilled engineers (including the hardware and software designers) than a typical client benchmark team.

Use benchmark results as a guide to the potential or *theoretical* performance and throughput of a server under a specific workload type and not as an indicator of *actual* performance.

2

IBM and benchmarks

Benchmarking is becoming more and more important to the IT industry. The IBM @server brand has had great success in industry standard benchmarks over the last couple of years, especially the xSeries with many #1 leading benchmarks.

This chapter explains the involvement of IBM and the IBM commitment to industry standard benchmarks. We discuss how the underlying design of the xSeries and the work of the IBM performance team made IBM #1 in benchmarks on the Intel® Xeon processor platform.

Following are the topics in this chapter:

2.1 Why IBM runs benchmarks on xSeries servers

Benchmarking is one of the few ways in the computer industry to objectively compare offerings between different computer vendors. Leading technology from IBM is compared against other server vendors in a standardized way of producing and publishing results.

Benchmarks are published records, and IBM invests much time, effort, and money to achieve the best possible results. All of the #1 results that xSeries solutions win are proof points that IBM technology made a huge difference in pushing the industry forward.

IBM takes benchmarking very seriously. This is primarily because our clients also take it seriously and consider the results an important differentiator. Additionally, IBM undertakes benchmarks as an integral part of the development process of the xSeries servers. The xSeries performance lab is actively involved in the process of bringing a new server to market to ensure that each system is properly tuned for its intended clients' use. We explain more about the work that these teams do in 2.4.1, "The xSeries performance lab" on page 10.

2.2 Client benefits

IBM makes a considerable investment in conducting industry standard benchmarks on xSeries servers. This is done to help clients make informed decisions, and to demonstrate technology leadership. However, the benefits go much deeper:

- Industry standard benchmarks provide a fair means of comparison. You can use the published results to compare and contrast systems, enabling you to make better hardware choices when architecting a solution.
- Benchmarking is an integral part of performance development. Use the data produced to evaluate a system's capabilities, improve designs, tune firmware, and highlight architectural strengths.
- The IBM Performance Lab works very closely with operating system, application, and database vendors. Analysis of server performance under these workloads can help identify issues and "hot spots" in the software. Performance improvements derived from this characterization are often distributed later in software updates and service packs, which benefit the industry as a whole.
- Servers are benchmarked with the same applications that clients run in production environments. Understanding these workloads enables IBM to help clients size servers and to assist tuning production environments.

2.3 Benchmarks important to IBM

The xSeries performance lab carefully selects which benchmarks to run and which xSeries servers on which to run them. Conducting industry benchmarks is a non-trivial exercise both in terms of effort and cost, so only the most important benchmarks are run. Benchmark plans are also carefully made well in advance, to ensure the best use of lab resources.

The reasons why specific benchmarks/server combinations are selected include the following:

- The benchmark represents a specific application type and the combination of that application and an xSeries server is a common client configuration. In general terms, a benchmark should be useful and relevant to clients.

  On the flip side, a benchmark that ceases to be relevant to clients may be dropped by xSeries.
- The industry considers specific benchmarks to be important for commercial workloads, and IBM wants to ensure that clients are able to compare xSeries servers with those from other vendors.
- The xSeries Performance Lab listens closely to client requests to produce specific benchmarks. When there is sufficient demand, the new requirement is put into the plan. If there is insufficient demand on some application benchmarks, the team may consider dropping a benchmark.
- IBM believes it can gain a leadership position in a specific benchmark on a newly announced server and feels clients benefit from such a benchmark result.

As a result, the xSeries performance lab concentrates on a subset of the available industry benchmarks. The list of industry benchmarks undertaken is very dynamic—benchmarks are retired or replaced with updated ones, which reflects the change in the way clients use their computing resources.

At the time of writing, following are the benchmarks for which the xSeries performance lab produces results:

System benchmarks

- TPC-C Benchmark (TPC-C) for transaction processing
- TPC-H Benchmark (TPC-H) for ad hoc decision support
- TPC-W Benchmark (TPC-W) for Web e-commerce
- TPC-App
- SPECWeb99 (Dynamic Web Content)
- SPECWeb99_SSL (Encrypted Dynamic Web Content)
- SPECcpu
- SPECjApp
- SPECjbb

Product-specific benchmarks

- Oracle OASB (Oracle Applications Standard Benchmark)
- PeopleSoft Enterprise
- PeopleSoft EnterpriseOne (former J.D. Edwards)
- Siebel PSPP (Platform Sizing and Performance Program) benchmark
- Citrix
- SAP R/3 2-tier SD benchmark
- SAP R/3 3-tier SD benchmark
- Microsoft Exchange MMB3 mail benchmark
- Lotus® Domino NotesBench mail benchmark

The development of benchmarks themselves is always evolving. There is an ongoing effort to meet the requirements of the industry; thus, benchmarks that are important now may be retired a year from now.

2.4 How xSeries fared

Since IBM announced Enterprise X-Architecture™ (EXA)-enabled server in 2001, the success was huge in producing #1 benchmarks. The EXA servers are the high-end family of xSeries, currently the x366, x445, and x455. xSeries EXA servers won over 80 number one standard industry benchmarks since they entered the market. This is not a coincidence—EXA made a huge impact in high performance.

There are three main reasons why xSeries won so many benchmarks.

- The xSeries performance lab
- The IBM @server family
- Enterprise X-Architecture

2.4.1 The xSeries performance lab

IBM puts significant effort into ensuring that its servers have the highest performance level possible. Part of this effort is the IBM xSeries Performance Lab, a group in Research Triangle Park, North Carolina where work is done on xSeries servers throughout the development phase and after they become available to clients.

During the development phase, performance models are created using subsystem and system functional specifications, chip functional specifications, and input from the IBM development engineering departments as well as trace information from the performance lab to do the following:

- Optimize the performance of the subsystem and system before the product is manufactured.
- Make design decision trade-offs.
- Select the optimum performance among various available chipsets that are intended for use as part of the subsystem or system.
- Select optimum settings of the chipset parameters.

This information is used to provide subsystem and system design guidance to the development engineering departments. As the system development phase nears completion, performance measurements are made with prototype subsystems, systems, as well as with ship-level systems to do the following:

- Perform stress testing.
- Validate product functional specifications.
- Validate the subsystem and system performance models.
- Optimize the performance of the subsystem and system.
- Develop performance white papers for marketing and demonstrating the competitiveness of the xSeries systems.
- Develop performance tuning guides for clients using specified applications.

Marketing, sales departments, and vendors can use this information to sell the xSeries systems. Clients can use this information to select the appropriate system and to tune their systems for their applications.

Figure 2-1 A TPC-C benchmark configuration in the xSeries performance lab

2.4.2 xSeries and HPC Benchmark Centers

The xSeries and high-performance computing (HPC) benchmark centers in Poughkeepsie, New York and in Montpeilier, France provide the capability of performing client benchmarking on Linux® and OpenPower™. This benchmarking includes proof of concept, scaling, and performance requests.

Skills in these centers focus on scientific and technical applications, which includes architecture porting, algorithms, compilers (Fortran 77, Fortran 90, C, C++), parallel programming and message passing, performance analysis and tuning, and job scheduling.

For more information, see the following Web sites:

- North America, South America, Asia-Pacific region (IBM employees only):

  http://benchmarks.pbm.ihost.com

- Europe, Middle East, Africa (EMEA) region (IBM employees only):

  http://w3.ibm.com/support/pssc

2.4.3 IBM eServer family

IBM is a high technology company and has been for many years. Server development has been one of the main areas of focus. In today's market, IBM has five strong server brands:

- IBM eServer zSeries®
- IBM eServer iSeries™
- IBM eServer pSeries®
- IBM eServer xSeries

- IBM @server BladeCenter™

All of these servers are high preforming servers that are state of the art in their part of the market. The great success of xSeries is partly dependent on the long experience that was gathered during many years of development and client feedback from the other server brands. Each brand is unique and fills special client requirements.

IBM @server zSeries - Zero downtime servers

The zSeries mainframe server was installed in enterprise clients for more than 30 years. The zSeries is built for environments with zero downtime requirements.

With zSeries, clients can support multiple business applications isolated from each other if required on a single server. zSeries can run computing environments like z/OS®, Linux, z/VM®, VSE/ESA™, IBM DB2®, WebSphere®, Lotus Domino, Tivoli® Storage Manager, and many more.

This server platform was once considered by many to be out of date and on its way out of the market, but there was a recent resurgence in the mainframe, partly due to its ability to run Linux. Many clients today run many instances of Linux with different types of applications on zSeries servers, and the reliability and uptime is undisputable in the server market.

IBM @server iSeries — integrated solution servers

This server platform is primarily used for integrated solutions. The iSeries server clients can run on computing environments for IBM i5/OS™ (the latest generation of IBM OS/400®), IBM AIX® 5L™, Microsoft Windows, Linux, IBM WebSphere, Lotus Domino, and Java solutions all on a single highly integrated and very powerful server.

In addition to running multiple operating systems simultaneously, iSeries servers support integrated xSeries servers to support Windows or Linux environments.

IBM @server pSeries — performance UNIX servers

pSeries server solutions offer the flexibility and availability to handle your most mission-critical and data intensive applications. They also deliver the performance and application versatility necessary to meet the dynamic requirements of today's infrastructure environments.

The pSeries servers are based on IBM Power 4 and Power 5 64-bit processors family and run on UNIX or Linux operating systems. When pSeries servers are designed high performance computing, scalability, and reliability are the lead words for the design team.

You can find solutions that leverage the pSeries server advantages everywhere from small departmental servers for computer-aided design (CAD) to large high performance clusters that are making large simulations.

AIX is the main operating system that runs on pSeries servers; however, Linux is becoming more and more popular.

IBM @server xSeries — X-Architecture Intel-based servers

xSeries is the Intel based server platform from IBM, and a fast growing one. This server family provides outstanding availability, scalability and price/performance capabilities that help better manage and provision IT infrastructures.

xSeries had great success in the market during the last couple of years above other server vendors. One of the main reasons for this huge success is the ability to take the best technologies from the other server platforms and incorporate it into xSeries servers. Each IBM @server brand contributed something to xSeries. For example:

- zSeries inheritance - RSA (Remote Supervisor Adapter)
- iSeries inheritance - Performance Manager in IBM Director
- pSeries inheritance - Ability to attach Remote I/O drawers to xSeries

Clients that run xSeries today can run all the common operating systems that are supported by the Intel x86 processor family like Windows, Linux, NetWare, ESX Server,

Clients that run xSeries today can run all the common applications that are supported by the Intel x86 processor family like, IBM WebSphere, Microsoft SQL server, Oracle, Java, Domino and many more.

> **Tip:** For a full compatibility list of xSeries hardware, applications, and middleware please visit the IBM ServerProven® web site:
>
> `http://www.pc.ibm.com/us/compat/indexsp.html`

IBM @server BladeCenter—where it all comes together

A couple of years ago many vendors introduced blade solutions to the computer market. The initial benefits of blade servers were to lower the total cost of ownership and to make administration easier.

The IBM offering, IBM @server BladeCenter is a highly modular chassis. It supports up to 14 hot-swappable, 2-way and 4-way Intel processor-based and 2-way POWER™ processor-based blades. It also supports a wide range of networking modules, including Gigabit Ethernet, InfiniBand, and Fibre Channel, for high-speed connectivity to the rest of your network. It also supports a redundant pair of management modules for complete systems management.

The BladeCenter is also an open server platform for other vendors of computer equipment to design their own solutions and integrate them into the BladeCenter chassis. Many large vendors like Nortel, Cisco, Brocade, and Qlogic released products that clients of BladeCenter can use today.

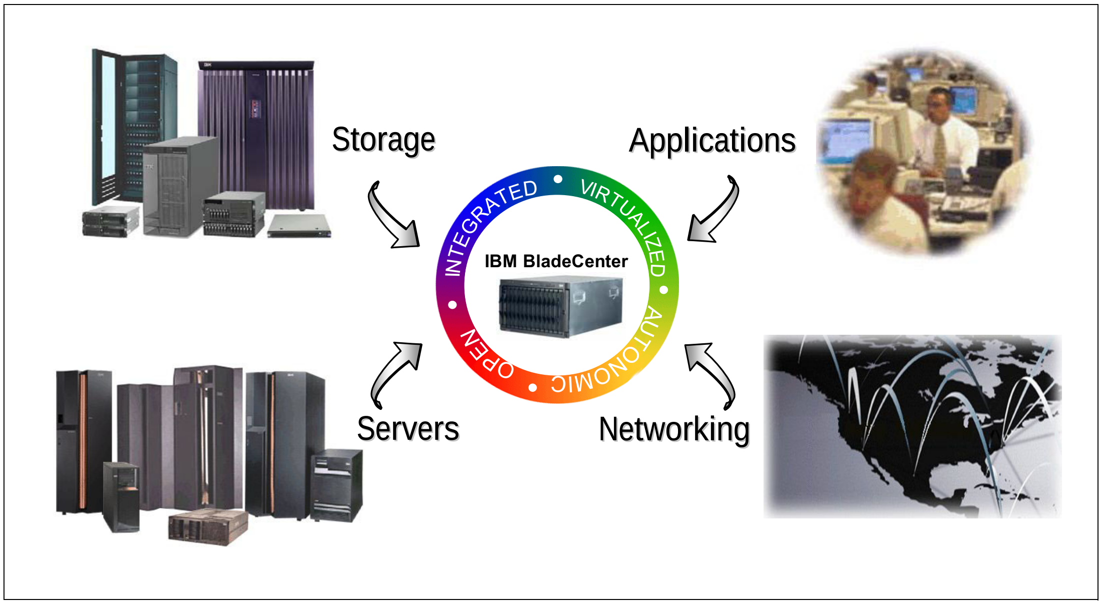

Figure 2-2 The BladeCenter philosophy

2.4.4 Enterprise X-Architecture technology

Let us look back a couple of years to 1998 when IBM introduced a technology blueprint called IBM X-Architecture. The idea for this blueprint was to extend and inherit the benefits of advanced mainframe technologies to Intel processor based servers. These benefits are in the areas of availability, scalability, systems management, service, and support.

Figure 2-3 The EXA inheritance from the IBM @server family

IBM continued this trend by extending the X-Architecture blueprint. The outcome became Enterprise X-Architecture technology. Most of this technology is delivered through IBM developed core logic. Core logic determines how the various parts of a system (microprocessors, system cache, main memory, I/O, etc.) interact.

2.5 How the industry benefits

There are two major areas that IBM contributes greatly with their work on benchmarking and performance.

- IBM is a member of the major benchmarking councils like TPC and SPEC.
- IBM together with other vendors of the computer industry are driving the development of benchmarks forward so that the benchmarks that are in the market today are useful for the clients.

We mentioned it earlier but it is worth repeating here, the work and the effort that the IBM performance lab does is not just to win benchmarks. Benchmarking is a part of the process of making the hardware and the software better. This is the second area where IBM contributes greatly and this area is more important from the clients point of view. This is also a major differentiator for IBM from other vendors.

When the xSeries performance lab sets up a solution for a benchmark it runs into bottlenecks very fast. The lab tunes the configuration and performs several changes per day to solve the bottlenecks and then they run the benchmark again. This is a process that is ongoing for months. The difference can be very big in performance.

The experience that the lab acquired during the test, by tuning and tweaking the configurations and all of the subsystems of a server, is then transferred to the development team and to the clients that buy the IBM xSeries servers. So the servers arriving to the clients are highly tuned.

But this is not the only area from which the clients benefit. When the performance lab does benchmarking it is not only dependent on xSeries server products. The operating systems and applications used on the xSeries servers are also bottlenecks and need to be tuned and developed. By doing this, the performance work performed helps the software vendors to develop their software as well.

This is done with all the major software vendors. We discuss two of the centers that IBM has.

- IBM Center for Microsoft Technologies
- Linux Technology Center

2.5.1 IBM Center for Microsoft Technologies

The IBM Center for Microsoft Technologies (CMT), located a few minutes from the Microsoft campus in Redmond, WA, is the primary interface IBM has with Microsoft. It supports products that run on all IBM platforms. IBM has highly trained IBM technical professionals at the CMT dedicated to exploiting Windows XP, Windows 2000, and Windows Server 2003 on Intel-based systems. The Center for Microsoft Technologies works in four areas:

- Developing device drivers for IBM products and Windows hardware abstraction layer (HAL) code for xSeries systems, including optimizing HAL code for best performance, and developing new technologies for the Windows platforms.
- Testing IBM systems in the IBM Microsoft-Certified Hardware Compatibility Lab for both the Microsoft-designed hardware compatibility testing (HCT) and the more demanding Microsoft system compatibility testing (SCT). IBM applications being developed for Windows operating systems are also tested for Microsoft standards compliance here.
- Defect Support with IBM Level 3 Support in high-severity situations when it is necessary to work directly with Microsoft Development personnel to resolve problems. The CMT also serves as a technical backup for the IBM Help Centers and as a worldwide center of IBM expertise in installation planning.
- Technical support for enterprise large accounts and independent software and hardware vendors developing products for IBM systems.

2.5.2 Linux Technology Center

The Linux Technology Center (LTC) serves as a center of technical competency for Linux both within IBM and externally. It provides technical guidance to internal software and hardware development teams and fulfills the role of an IBM extension to the open source Linux development community.

The LTC is a world-wide development team in IBM whose goal is to use the IBM world-class programming resources and software technology to actively accelerate the growth of Linux as an enterprise operating system while simultaneously helping IBM brands exploit Linux for market growth.

The LTC currently has programmers involved in many Linux projects including scalability, serviceability, OS security, network security, networking, file systems, volume management, performance, directory services, standards, documentation, accessibility, test, security certification, systems management, cluster management, high availability, storage & I/O, PowerPC® support, power management, reliability, internationalization, and other projects required to make Linux a mature operating system ready for mission critical workloads.

Members of the LTC work directly in the open source community using standard open source development methodology. They work as peers within the shared vision of the Linux community leadership and participate in setting Linux design and development direction.

2.6 Examples of xSeries benchmark setups

Performing a benchmark is not a trivial task. It requires time, hardware, effort and the cost can be significant, especially for hardware-intensive benchmarks such as TPC-C.

In this section we describe two benchmark setups performed by the xSeries Performance Lab. One of them is a large setup, the other is smaller. We describe the hardware requirements to do industry standard benchmarks.

2.6.1 TPC-C on an xSeries x445

In this example we took one of the largest industry standard benchmarks as an example, TPC-C.

The benchmark team used a 16-way x445 machine running Windows Server 2003, Datacenter Edition as an operating system and SQL Server 2000, Enterprise Edition as a database engine. For client, eight x225 servers where used to emulate 172,000 users working with the database.

The largest part of the hardware configuration is the disk subsystem. It contained six IBM TotalStorage® DS4400 controllers (FAStT700), 44 IBM TotalStorage EXP700 disk enclosures with a total of 616 hard drives (14 each). The exact configuration and all of the components used appear in Figure 2-4 on page 17.

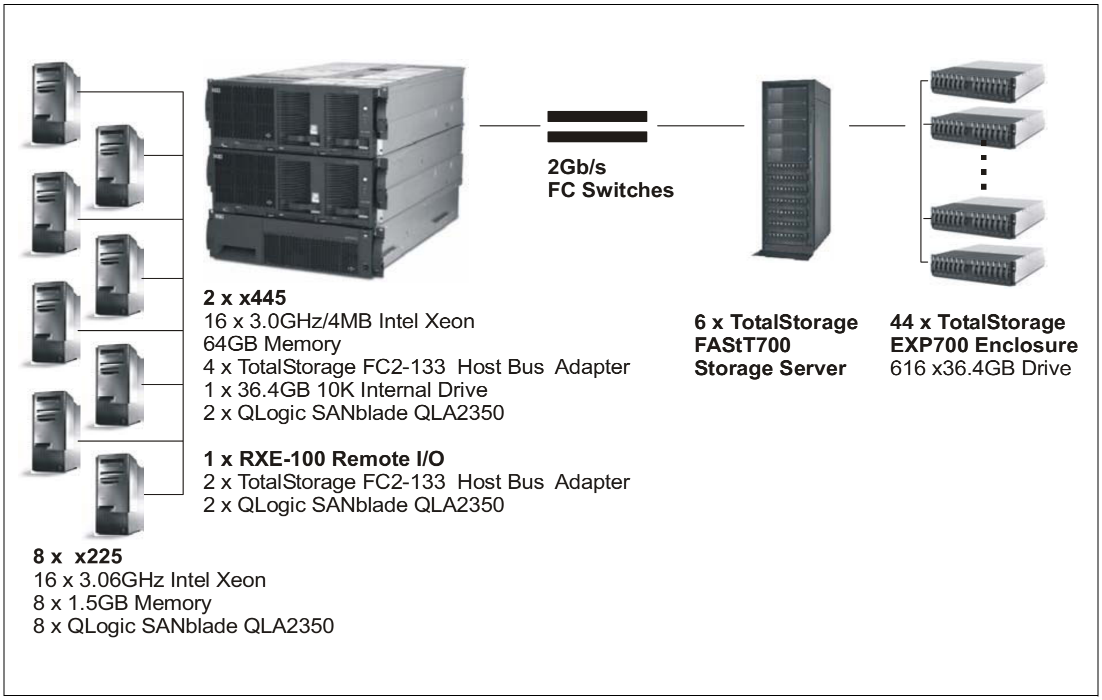

Figure 2-4 TPC-C benchmark configuration

For TPC-C, the bottleneck is the disk subsystem due to the benchmark having a very high random I/O workload, which is typical of Online Transaction Processing (OLTP) applications; hence, the need for a large number of disk drives—needed to run the 16 processors to maximum utilization. To further enhance disk throughput, each drive only contained a small amount of data and the data was carefully positioned on the disks to maximize throughput and to minimize disk head movement.

SQL Server was configured to use almost all of the available 64 GB of system memory as a single buffer pool. A small amount of memory was reserved for the operating system and I/O device overhead. The SQL Server processor affinity mask was set to use the 16 physical processors in the x445 and the Hyper-Threading component of the processors. SQL Server distributed work across 32 processor instances. The operating system was configured to address more than 4 GB of memory by adding the /PAE switch to the Windows boot.ini file.

2.6.2 SPECweb99 on BladeCenter

From a hardware point of view, the SPECweb99 benchmark is much smaller than the TPC-C benchmark requirements. Even so, the configuration described here is unusually large for a SPECweb benchmark.

In this configuration, the EM64T-based HS20 blade servers are running SUSE LINUX Enterprise Server 9, 64-bit edition. Zeus is the Web server application used.

The disk subsystem is a DS4400 (FAStT700) with two EXP700 expansion units fully loaded with 73 GB hard drives. The clients are simulated using 10 x 226 servers that emulate Web users making requests to the Web servers.

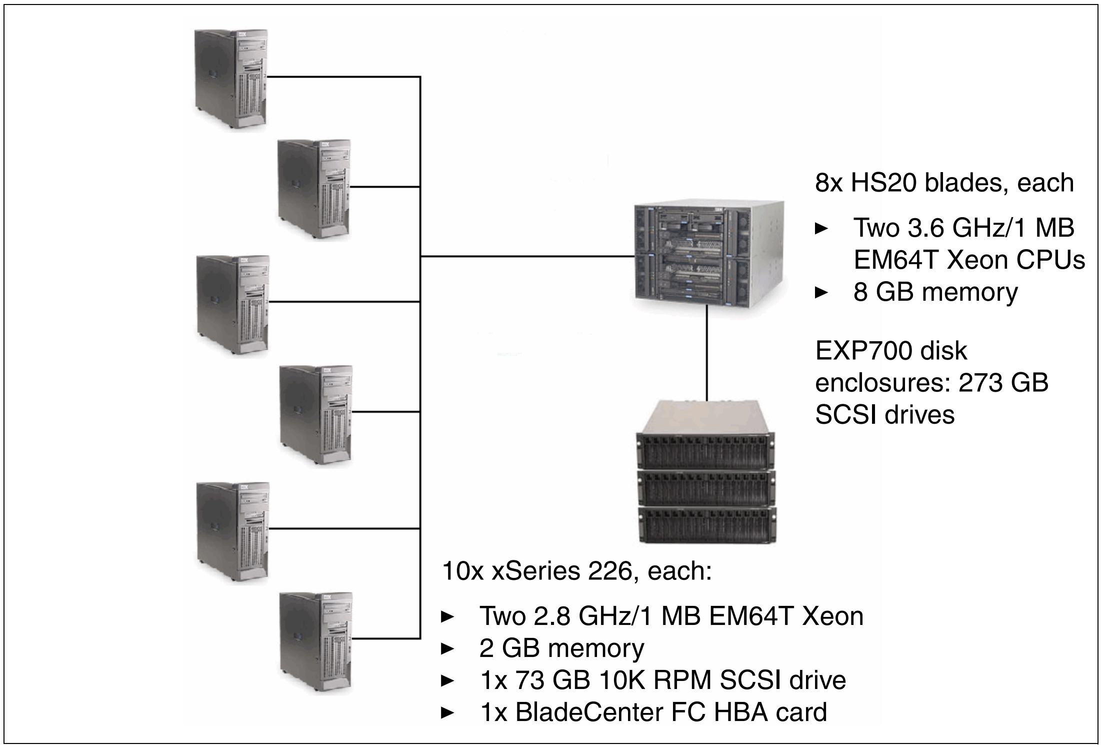

Figure 2-5 Real life setup made by IBM for a SPECweb99 benchmark

3

Industry benchmarks

Industry standard benchmarks evolved from a systemic problem within the industry, where every vendor claimed to offer the best value and to have the fastest system. Unfortunately the information provided was usually based on the vendor's own criteria. A level playing field was needed to provide clients with standardized performance information to assist in determining the solution that best suits them.

Industry standard benchmarks are a valuable resource, that clients can use as an aid when making decisions regarding hardware and software platforms. There is no substitute for testing a system using the intended workload, but if this is not possible, clients can use the data from industry standard benchmarks as a part of their evaluation process. References, capacity planning, total cost of ownership and performance modeling, should also be a part of this process.

Following are the topics in this chapter:

3.1 What benchmarks are available

Industry standard benchmarks provide a fair means of comparing system performance. You can use the published results to compare and contrast systems, enabling clients to make better hardware choices when they architect a solution.

The TPC and SPEC maintain the most widely recognized and accepted industry standard benchmarks. Both of these entities are non-profit organizations formed to create, maintain, and monitor effective performance benchmarks. There are also many benchmarks that are specific to an application, for example Oracle Applications Standard Benchmark, Microsoft Exchange MAPI Messaging, and SAP Standard Application Benchmark. All of these benchmarks are referred to as industry standard because they are regulated in some way to ensure that the published results are credible.

The following sections describe who the industry standard bodies are, and how the benchmarks are regulated.

Don't try this at home: Industry standard benchmarks can be very complex and expensive to perform. In most cases it is simply not practical for clients to perform the benchmarks themselves. Verified results for most xSeries servers are freely available from TPC, SPEC, or other standards bodies.

3.1.1 Transaction Processing Performance Council

The Transaction Processing Performance Council (TPC) was founded in 1988 to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry.

The TPC benchmarks simulate complete transactional computing environments, replicating large multi-tasking/multi-user workloads. In contrast to most other benchmarks, TPC benchmarks are modeled after actual production applications and environments rather than stand-alone computer tests that may not evaluate key performance factors like user interface, communications, disk I/O, data storage, and backup and recovery.

The requirements are described at a functional level, rather than distributed via source code. It is then up to the test sponsor, which is usually a collaboration between the hardware and database vendors, to build the environment and perform the benchmark. A full disclosure report must be submitted as proof that all of the requirements were met, and the result is verified by an independent auditor before it can be published.

The TPC benchmarks tend to be the most difficult and expensive benchmarks to run. The benchmark requirements are described at a functional level, so software harnesses need to be developed to generate the load, and the environments are often large, consisting of multiple servers and clients. The benefit of this approach is that the systems are stressed and tested as a whole, which more closely approximates workloads seen in real environments.

3.1.2 Standard Performance Evaluation Corporation

The Standard Performance Evaluation Corporation was established in 1988 by a small group of workstation vendors who realized that the marketplace was in need of realistic, standardized performance tests. The idea was that an ounce of honest data was worth more than a pound of marketing hype.

Originally formed to create and monitor standardized CPU benchmarks, SPEC evolved into an umbrella organization encompassing three diverse groups:

- The Open Systems Group (OSG)

  The OSG is the original SPEC committee. This group focuses on benchmarks for desktop systems, high-end workstations, and servers running open system environments. The benchmarks include, CPU, Java, mail, and web serving.

- The High-Performance Group (HPG)

  These benchmarks target high-performance system architectures, such as symmetric multiprocessor systems, workstation clusters, distributed memory parallel systems, and traditional vector and vector parallel supercomputers.

- The Graphics Performance Characterization Group (GPC)

  The SPEC/GPC Group is the umbrella organization for project groups that develop consistent and repeatable graphics benchmarks and performance-reporting procedures. SPEC/GPC benchmarks are worldwide standards for evaluating performance in a way that reflects user experiences with popular graphics applications. These benchmarks are focused at workstations rather than servers.

The basic SPEC methodology is to provide the vendor with a standardized suite of source code based upon existing applications that were already ported to a wide variety of platforms by its membership. The vendor then takes this source code, compiles it for the system in question, and then tunes the system for the best results. The use of already accepted and ported source code greatly reduces the problem of making apples-to-oranges comparisons.

The SPEC benchmarks are generally component focused. These benchmarks are very good for isolating individual components for direct comparison, but take care when using them to reach meaningful conclusions, as the interaction system elements not tested often impact performance in real application workloads.

3.1.3 Application and other standardized benchmarks

There are a number of application specific benchmarks available, such as SAP Standard Application, Oracle Application Standard, Notesbench, and Microsoft MAPI Exchange. How these benchmarks are developed and maintained varies, but they are all regulated in some way to ensure that the published results are fair. Some require independent auditors, while other have vendor neutral teams within the application vendors organization that verify the results.

These benchmarks are used widely to compare servers and operating systems once a decision on a particular application is made. They are often also used as rough sizing guides, although great care must be taken when doing this.

3.2 Rules of engagement

Standardized benchmarks were developed to enable systems to be compared. An integral part of this process is a mechanism to make sure that the rules of a benchmark are followed. This section briefly discusses the methods that TPC, SPEC, and a few of the other key benchmark providers use to regulate their benchmarks.

3.2.1 TPC rules of engagement

The rules of TPC benchmarks are strictly enforced. Before publication, an independent auditor must review and verify a benchmark result to ensure that all of the requirements were met. A full disclosure report is filed with the TPC administrator detailing a component's, hardware, software, and all parameter settings required to reproduce the result.

The full disclosure report also gives competitors the opportunity to learn tuning techniques used to achieve the result, and to challenge the result if they believe any of the rules were broken. A Technical Advisory Board reviews challenges to a benchmark result. If this board decides that the challenge has merit, then the members of the TPC vote on the board's recommendations, and may withdraw the result.

The rules and process to enforce them lends credibility to the results. To further build on the credibility, TPC promotes the concept of "Fair Use" for its members when using TPC related information. These include:

- All of the primary metrics must be reported for a benchmark.
- A full disclosure report must be complete and on file with the TPC administrator.
- Members cannot use estimated results that refer to the TPC or TPC workloads, or compare them to results.
- Members cannot report one primary metric without including the other.
- Members cannot make TPC-related claims or lead the reader to TPC-related conclusions that are untrue or cannot be substantiated by the entire body of results.
- Members cannot compare one system's total price to the partial price of another system, or compare partial price to partial price. Results cannot be generated with less than the entirety of the configured system.
- Members cannot refer to a withdrawn result without specifically stating that the result is withdrawn.

3.2.2 SPEC rules of engagement

This section addresses the guidelines for submitting benchmark results for review and publication on the SPEC Web site.

Each benchmark suite produced by the SPEC includes a set of specific run and reporting rules that must be followed to produce a publishable result. These rules require the following:

- Proper use of the SPEC benchmark tools as provided
- Availability of an appropriate full disclosure report
- Support for all of the appropriate protocols and standards

Additionally, SPEC expects that any public use of results from this benchmark suite are for systems and configurations that are appropriate for public consumption and comparison. This means the following:

- Hardware and software used to run this benchmark must provide a suitable environment for the applications the benchmark targets.
- Optimizations used must improve performance for a larger class of workloads than just the ones that this benchmark suite defines.
- The tested system and configuration is generally available, documented, supported, and encouraged by the providing vendors.

Like TPC, SPEC fosters the concept of “Fair Use”. There are a number of rules around reporting benchmark results. This is to build credibility and to prevent misuse. These rules include the following items:

- The basis for comparison must be stated.
- The source of the competitive data must be stated.
- The date competitive data was retrieved must be stated.
- All data used in comparisons must be publicly available (from SPEC or elsewhere).
- The benchmark must be currently accepting new submissions if previously unpublished results are used in the comparison.

3.3 Selecting the relevant benchmark

The published results of industry standard benchmarks are an extremely useful approximate guide to relative performance. The most important aspect of examining these results is to find the benchmarks that best represent the intended workload.

You can divide most commercial workloads into one of the following categories:

- Online Transaction Processing (OLTP)
- Decision Support (DSS)
- Application servers
- Enterprise Resource Planning (ERP)
- Web Services
- Web serving
- Mail and collaboration

If a client knows what kind of workload their environment has, for which benchmark do they look? We summarized this relationship in Table 3-1 that shows different types of workloads and the benchmarks that apply to them. We discuss some of the workloads in greater detail in 3.4, “Workload descriptions” on page 24 and some of the benchmarks in 3.5, “Benchmark descriptions” on page 27.

Table 3-1 Different workloads and industry standard benchmarks that apply to them

Application type	Relevant benchmarks
OLTP	TPC-C
DSS	TPC-H
ERP: PeopleSoft	PeopleSoft Enterprise
ERP: PeopleSoft EnterpriseOne	(Formerly J.D. Edwards) PeopleSoft EnterpriseOne 3-Tier
ERP: Baan	BaanERP
ERP: SAP R/3	SAP SD 2-Tier and SAP SD 3-Tier
ERP: Siebel	Siebel PSPP (Platform Sizing and Performance Program)
Oracle	OASB (Oracle Applications Standard Benchmark)
Java Server	SPECjApp Server
Java VM Server	SPECjbb2000
Java VM Client	SPECjvm98

Application type	Relevant benchmarks
Web Services	TPC-App
Web Servers	SPECWeb99 and SPECWeb99_SSL
Mail: Exchange	SPECmail2001 and Microsoft Exchange MMB3
Mail: Domino	SPECmail2001 and NotesBench
Terminal Server/Citrix	CSTK (Citrix Server Test Kit), LoadRunner, Microsoft Terminal Server capacity tools
CPU-intensive workloads	SPEC CPU2000
HPC workloads	Linpack, SPEC HPC, Palas MPI

3.4 Workload descriptions

As we mentioned in 3.3, "Selecting the relevant benchmark" on page 23, most commercial workloads can be divided into a few major workload categories. In this section, we describe what every category does.

3.4.1 Online Transaction Processing (OLTP)

Online Transaction Processing (OLTP) systems are generally used by a large number of concurrent users executing transactions against a centralized database. The transactions are usually simple, but require immediate update and response from the database. OLTP workloads can vary widely, but are usually characterized by a high number of random I/O requests for the database servers. OLTP workloads are often very I/O intensive due to simple transactions requiring random disk reads and disk writes. Potential areas that have the most impact on performance are:

- Memory subsystem

  Buffer caches are one of the most important components in the server. If the server does not have sufficient memory paging occurs, resulting in excessive disk I/O, increased disk latencies, and increased transaction response time for the end user.

- Disk subsystem

  Even with sufficient memory, most database servers perform large amounts of disk I/O to bring data records into memory and to flush modified data to disk. It is important to configure a sufficient number of disk drives to match the CPU processing capability. With many OLTP database applications, performance is limited by the large number of random I/O requests. In this situation, adding additional drives or disk controllers can help to improve performance.

- CPU subsystem

  Processing power is another important factor for database servers as database queries and update operations require intensive CPU time. The database replication process also requires considerable amounts of CPU cycles. Database servers are multi-threaded applications, so SMP-capable systems provide improved performance scaling to 16-way and beyond. Level 2 (L2) cache size is also important due to the high hit ratio. For example, SQL Server's L2 cache hit ratio approaches 90%.

- Network subsystem

  The networking subsystem tends to be the least important component on an application or database server because the amount of data returned to the client is a small subset of the total database. The network can be important, however, if the application and the database are on separate servers.

3.4.2 Decision Support (DSS)

Organizational decision-makers use Decision Support Systems to improve strategic, tactical, and operational decisions. Like OLTP environments, Decision Support environments center around a relational database, but generally there are fewer concurrent users, and complex queries are much more prevalent. System throughput is key to a DSS environment as the workloads tend to be very sequential in nature, performing a large number of sequential accesses to memory and disk because of full table and index scans.

An important characteristic of DSS workloads is the relatively low buffer and cache hit rates when compared to OLTP workloads. This is due to a combination of the large data sets, and sequential nature of accesses.

- Disk subsystem

  The I/O subsystem is very important in a DSS support system, and should be maximized for throughput. However, disk access tends to be very sequential, so a smaller number of drives are required than in OLTP environments, where many spindles are needed to reduce the random access latency. The database logs are also accessed less frequently as data is modified less often.

- CPU subsystem

  Processor speed is important for DSS environments, but the size of the level 2 (L2) cache is not. This is because of the low cache hit rate characteristic of DSS environments.

- Memory subsystem

  The complex queries that are more prevalent in DSS workloads use memory intensive operators like sort and hash-join; therefore, a system needs to be configured with a sufficient amount of memory for each processor. However, because of the low buffer hit rate typical of DSS workloads, adding more and more memory results in only small performance improvements.

3.4.3 Web servers

Web server deployments range from machines serving only a small set of static files to e-commerce solutions requiring database driven dynamic pages and distributed transactional applications. The workload of a Web server depends on the content it serves. For predominantly static content, the speed of the network connection between the server and users may be the limiting factor. For dynamic and mixed content, the memory subsystem is the major factor in performance. Finally for secure content, the CPU heavily influences system performance.

- Memory subsystem

  A Web server requires memory to run applications, cache the requested data, and to generate dynamic content. This is generally the most significant subsystem affecting performance.

- CPU subsystem

  The CPU subsystem can be very heavily utilized if a server is using SSL to encrypt the traffic. A considerable amount of CPU processing is required to encrypt and decrypt data.

- Network

  The network can limit system performance for a Web server, but this is usually due to bandwidth limitations between the server and client, rather than the network interface of the server itself.

- Disk subsystem

  The disk subsystem is usually the least concern for a Web server. With enough memory installed, a Web server caches most of the content in memory, greatly reducing disk access.

3.4.4 Application server and ERP

Application and ERP based workloads are emulating real user interaction with the systems in one or more application modules of the benchmarked application suite. Normally an application or an ERP system vendor uses the module that their client most commonly uses.

One very interesting thing about application and ERP solutions is that they are designed in two ways, as a 2-tier or a 3-tier solution. A 3-tier solutions consists of the following:

- Front end layer
- Middleware application layer
- Back-end database layer

In a 2-tier solution, the middleware and the back-end are integrated as one part. In most cases applications and ERP solutions are designed as 3-tier solutions.

We have different characters on the workload depending on the design. Generally speaking the middleware is the layer that is normally most constrained in a 3-tier solution because that is where all the applications are.

For more details on the different workloads that application and ERP vendors use, refer to their Web sites. SAP and Oracle's are as follows:

- SAP: `http://www.sap.com/benchmark/`
- Oracle: `http://www.oracle.com/apps_benchmark`

3.4.5 Mail and collaboration

Messaging and collaboration systems have become an integral part of business functions. How they are used and the workloads that they impose on a system varies markedly with different implementations. On a basic level, e-mail servers act as repositories and routers of electronic mail, but they can also act as Web applications and database servers. In general the important subsystems are as follows:

- Disk subsystem

  Mail and collaboration servers are often very disk intensive. Paging files, particularly for Lotus Domino, can be used extensively and should be separated from data. Log files are also very heavily used, and should be separated from data.

- Processor subsystem

  Many messaging and collaboration functions are CPU intensive, such as routing mail and indexing and searching databases.

- Memory subsystem

  It is very important that the mail servers have sufficient memory to avoid excessive paging. However, most mail and collaboration applications are 32-bit and do not make use of memory above 2 GB.

3.5 Benchmark descriptions

In this section we describe some of the major industry standard benchmarks used by the industry today. We cover a general overview of the benchmark, why it is important for the industry, and some constraints that the different benchmarks have.

3.5.1 TPC-C

The Transaction Processing Council's TPC-C benchmark is the industry's most often quoted benchmark. It is the third in a series of benchmarks that measure performance and price of online transaction processing systems. It replaced the previous TPC-A and TPC-B benchmarks. It simulates a complete order entry computing environment for a wholesale supplier, simulating operators executing transactions against a database. Transactions include order entry, delivery, recording payments, checking order status, and monitoring stocking levels.

Three metrics must be reported for this benchmark, the performance of the system, total cost, and the date that the system is available in its entirety. The reported performance is the number of complete orders processed per minute, measured in transactions per minute (tpmC). The total system cost includes hardware, software, three years maintenance, and enough storage capacity to accommodate 60 eight-hour days of operation.

For more details of the TPC-C benchmark, visit the following Web site:

http://www.tpc.org/tpcc/

Why is TPC-C important

The degree of interest in the TPC-C benchmark reflects the nature of the systems that it represents. Online transaction processing is the core computing system for many businesses. If it becomes overloaded, or fails, business functions may halt with far reaching consequences. The TPC-C benchmark tests the system in a large multi-user environment, and simulates many transactions that are common in real OLTP systems.

For hardware vendors there are additional benefits to the TPC-C benchmark.

- TPC-C stresses the system enormously: the CPUs, memory, I/O subsystems, and the database. There are many instances where defects, or constraints in hardware, firmware, operating system, and databases were uncovered from TPC-C benchmark analysis.
- The TPC-C benchmark produces repeatable performance results if the system configuration is unchanged. The benchmark workload and system performance reach a steady state, where the system and subsystem performance can be analyzed, and the effects of performance tuning changes can be easily measured.
- The TPC-C benchmark is one of the most highly characterized workloads. Much of the information regarding how processor architecture changes effect system performance is obtained from studying this benchmark.

Characteristics and constraints

Figure 3-1 shows a typical client/server configuration for a TPC-C benchmark. The database server in this example is an 8-way xSeries 445, configured with 64 GB of memory and in excess of 400 disks. A proprietary Remote Terminal Emulator (RTE) program, simulating 124,800 terminal users, executes transactions against the database. Four xSeries 225 servers were used as benchmark clients, handling screen I/O for the terminals and requests to the database server.

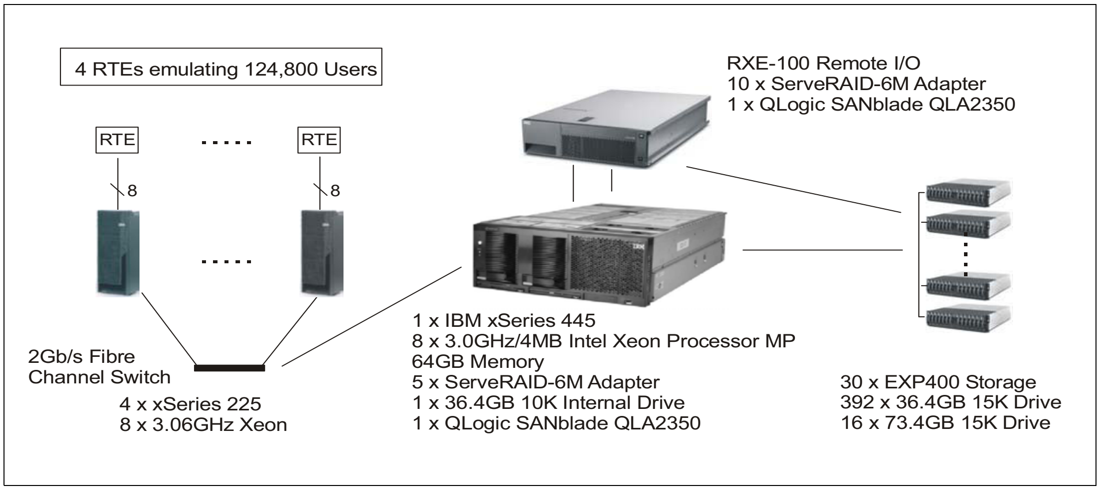

Figure 3-1 Typical configuration for a TPC-C benchmark

The striking aspect of this configuration is the large number of hard drives used. The performance of a server in a TPC-C benchmark is heavily influenced by random I/O latency. In order to move the bottleneck from the I/O subsystem to the CPUs, a large number of drives are added to the system, providing more disk heads to handle the random I/O of the workload.

Note: TPC-C can be a very expensive benchmark to run, due to the large numbers of disks required. There is another benchmark currently in development, TPC-E, that is likely to become an alternative for OLTP-type workloads.

3.5.2 TPC-H

TPC-H is a decision support benchmark that illustrates decision support systems (DSS) that examines large volumes of data, executes queries with a high degree of complexity, and gives answers to critical business questions. You cannot use pre-existing knowledge of the queries to optimize the system, so query execution times can be very long.

The performance metric reported by TPC-H is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size), which reflects multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries are executed, the query processing power when queries are submitted by a single stream, and the query throughput when queries are submitted by multiple concurrent users. The TPC-H price/performance metric is expressed as $/QphH@Size.

Results are published on the TPC Web site according to the size of the database. Currently there are 100 GB, 300 GB, 1,000 GB, 3,000 GB, and 10,000 GB categories.

For more details of the TPC-H benchmark, visit the TPC Web site:

http://www.tpc.org/tpch/

Why is TPC-H important

Businesses collect enormous amounts of data every day: information regarding orders, inventory, accounts payable, point-of-sale transactions, and clients. The amount of data increases exponentially, on average doubling every two-to-three years. Consolidating and organizing data for better business decisions can lead to a competitive advantage. Learning to uncover and leverage those advantages is what Business Intelligence is all about.

Business Intelligence often leverages information and data rules engines to help make these decisions along with statistical analysis tools and data mining tools. Whatever the tools or methods used, ultimately queries are run against a relational database management system (RDBMS) to find the relevant data for analysis. This back-end system is known as a decision support system (DSS), or data warehouse, and has become a cornerstone of many successful enterprises. It supports not only traditional strategic decision-making, but also day-to-day operational business processes. The decision support system has become a critical resource for large numbers of users, supporting a wide variety of both simple and complex processes.

Characteristics and constraints

TPC-H is a server only benchmark that has no clients or network component. A typical non-clustered configuration is shown in Figure 3-2. The workload is highly sequential, so systems are optimized for throughput. There are elements of random access to the temporary table spaces during some operations, such as table-joins, but predominantly throughput is the influencing factor effecting system performance.

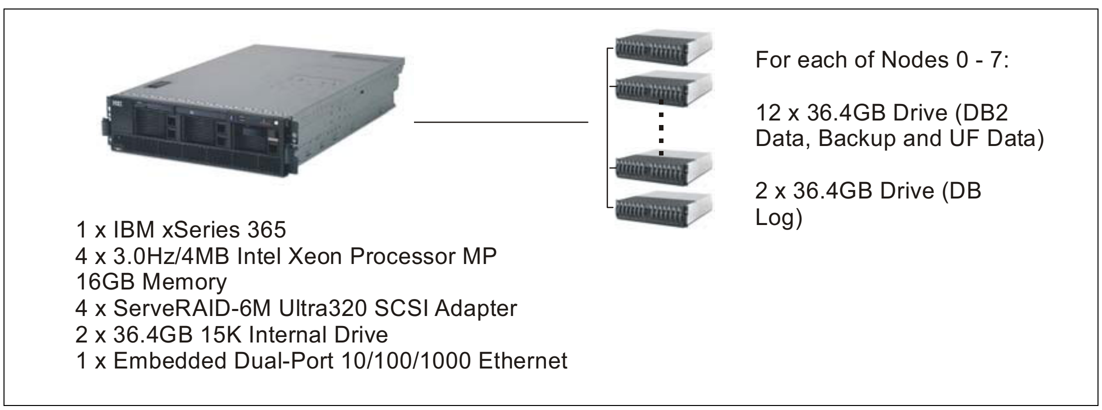

Figure 3-2 Typical non-clustered TPC-H configuration

The TPC-H workload has low cache hit rates, so increasing memory size has much less effect on performance than TPC-C. Also, in contrast to TPC-C, the system does not achieve a steady-state during the TPC-H benchmark. CPU utilization peaks and troughs throughout the benchmark as the various queries are submitted.

Another characteristic of the TPC-H is that the workload is highly partitionable. This means that the database can be split amongst separate clustered servers, with a very good scaling factor. This is ideal for highlighting the scale-out approach for databases, but it does mean that very high performance figures can be achieved for the benchmark by adding more and more servers.

3.5.3 TPC-W and TPC-App

TPC-W is a transactional Web benchmark, modelling an e-commerce Web storefront with clients browsing and buying books from a retail store Web site. TPC-App will soon replace it due to limitations in the benchmark. TPC-App was originally named TPC-W Version 2.

In this Redpaper, we discuss TPC-App as though vendors have not been submitting TPC-W benchmarks for some time. This is largely because the application was a custom written C++ application, rather than a managed environment such as Java, or .NET, and the specification permitted large portions of the workload to be cached, resulting in massive scale-out systems of 50+ servers.

TPC-App was designed to benchmark a commercial application server in a business-to-business Web services environment. Key differences from TPC-W are shown in Table 3-2.

Table 3-2 Key differences between TPC-W and TPC-App

TPC-W	TPC-App
E-Commerce Web Storefront	B2B Web Services Workload
Optimized hand written static code (C++)	100% managed - for example Java or .NET
Relaxed on caching and offloading opportunities	Tight on caching and offloading opportunities
~30% SSL	100% SSL (Secure, encrypted communications)
Showcase for image caches, Web caches, Web servers	Showcase for commercial application server environment

The basic operations simulate the business activity of a distributor supporting user online ordering and browsing activity. The application accepts incoming Web service requests from other businesses (or a store front) to place orders, view catalog items and make changes to the catalog, update or add client information, or request the status of an existing order. The majority of the request activity is to generate order purchase activity with a smaller portion of requests browsing the site. The workload is centered on business logic involved with processing orders and retrieving product catalog items and provides a logical database design.

The performance metric reported by TPC-App is the number of Web services processed per second. Multiple Web service interactions simulate the business activity, and each interaction is subject to a response time constraint. The performance metric for this benchmark is expressed in Web Service Interactions Per Second (SIPS).

All references to TPC-App results must include the primary metrics: the SIPS rate, the associated price per SIPS ($/SIPS), and the availability date of the priced configuration.

Why is TPC-App important

Web services describes a standardized way of integrating Web-based applications using the XML, SOAP, WSDL, and UDDI open standards over an Internet protocol backbone. XML tags the data, SOAP transfers the data, WSDL describes the services available, and UDDI lists the available services. Used primarily as a means for businesses to communicate with each other and with clients, Web services allows organizations to communicate data without intimate knowledge of the others' IT systems behind the firewall. This solves one of the most critical problems that businesses face.

Web services are used in a range of application integration scenarios—from simple, ad hoc, behind-the-firewall, data sharing to very large-scale Internet retailing and currency trading.

TPC-App is intended to benchmark an application server in a realistic business-to-business Web services environment.

Characteristics and constraints

Although the benchmark simulates a complete environment, the focus of TPC-App is the application server. Although originally designed to test only a single application server, the benchmark now allows the application layer to be scaled out across multiple servers (cluster). This was made because it is representative of real world environments. However, the workload is designed to try and limit this scaling so massive that configurations, such as those used in TPC-W, do not occur. Figure 3-3 shows a sample TPC-App configuration.

At the time of writing this Redpaper, the TPC-App specification was not yet accepted, so information regarding the benchmark and workload characteristics is limited.

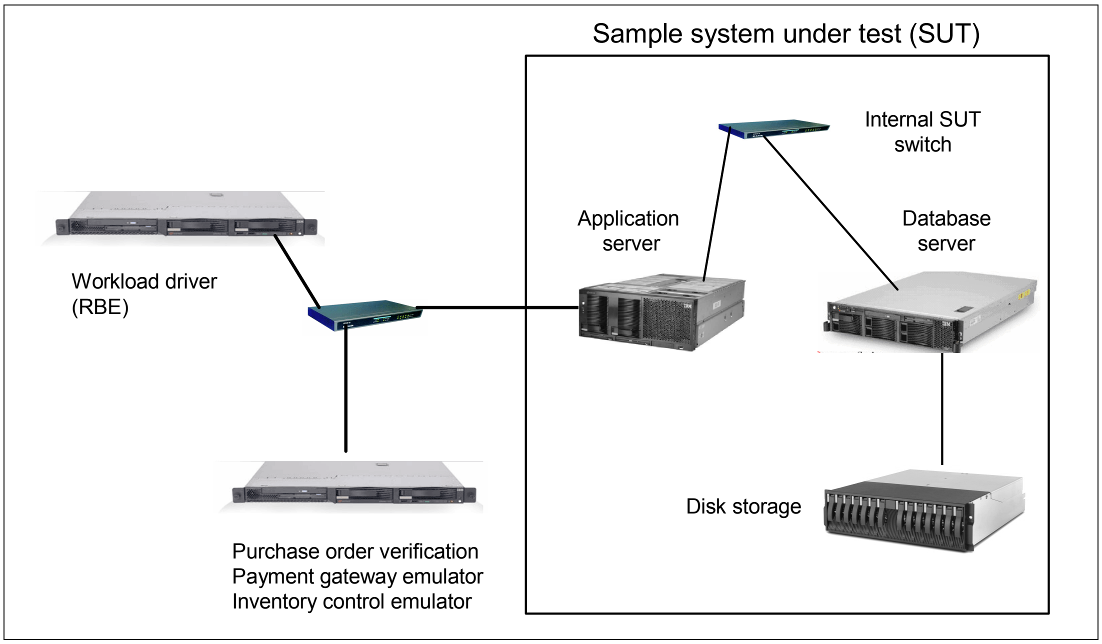

Figure 3-3 Sample TPC-App configuration

3.5.4 SAP Standard Application Benchmark

As defined on SAP's Web site, SAP Standard Application Benchmarks test and prove the scalability of mySAP Business Suite. The benchmark results provide basic sizing recommendations for clients by testing new hardware, system software components, and Relational Database Management Systems (RDBMS). They also allow for comparison of different system configurations.

The benchmarking procedure is standardized and well defined. The SAP Benchmark Council, made up of representatives of SAP and technology partners involved in benchmarking, monitors it. Originally introduced to strengthen quality assurance, the SAP Standard Application Benchmarks also tests and verifies scalability, concurrency and multi-user behavior of system software components, RDBMS, and business applications. All performance data relevant to system, user, and business applications are monitored during a benchmark run and can be used to compare platforms and as basic input for sizing recommendations.

There are several benchmarks available from SAP, but the most commonly performed and compared is the Sales and Distribution (SD) benchmark. This benchmark uses the Sales and Distribution model, which is very system intensive, and has 2-tier and 3-tier categories.

- SD 2-Tier benchmarks a system with both the database and the SAP application code on the same server.
- SD 3-Tier allows the SAP application to be separated out to a middle layer of application servers.

Why is SAP Standard Application Benchmark important

The implementation of SAP R/3 systems is a complex and expensive effort that has a major impact on an organization. Since ERP systems usually support critical internal business processes, an SAP implementation requires investment of IT resources to guarantee the achievement of system operational requirements, particularly performance, availability, and serviceability.

Additionally processor vendors like to use the SAP SD benchmark to demonstrate technological advances. It is a very CPU intensive benchmark that highlights the effect of L2 cache, and allows very good scalability in performance.

Characteristics and constraints

The SAP SD application benchmark is one of the most CPU intensive tests certified by SAP, and most commonly used for benchmarking. There is a 2-tier benchmark where the database and SAP application reside on the same server, and a 3-tier benchmark where the SAP application is separated out into a middle layer of application servers.

In the 2-tier environment, when the system reaches 100% utilization, the SAP application consumes about 80% of the CPU, and the database only approximately 8-10%. The performance measure is related to the number of users the system can sustain with a response time of less than two seconds.

The 3-tier SAP SD benchmark allows the SAP application layer to be scaled-out across multiple application servers that are connected to the database server. This is a very scalable environment, as the system can support more and more users by adding additional application servers until the database server reaches saturation. However, because the SAP application accounts for such a large proportion of the workload, when separated out a large number of application servers can be supported by the database server. Additionally, by increasing the size of the database server, even more application servers can be added, and even more users supported.

Deep pocket benchmark: The SAP SD 3-tier benchmark is sometimes referred to as a *deep pocket* benchmark. This means that the performance metric (number of benchmark users) can be increased by adding more and more application servers and increasing the size of the back-end database. You can obtain very high performance figures by simply spending more money, resulting in configurations of 100+ application servers.

Rather than examining systems that obtained the maximum performance figure for this benchmark, examine systems that are of a reasonable configuration for the intended environment. For example if 5, 10, or 20 application servers are reasonable for the intended workload, then examine the performance of benchmarked systems of approximately that size.

3.5.5 Oracle Application Standard Benchmark

Oracle Applications Standard Benchmark (OASB) is a benchmark of Oracle applications. Although there is a database component, this is not a database oriented benchmark.

OASB is a comparable standard workload that demonstrates the performance and scalability of Oracle applications and provides metrics for the comparison of Oracle applications' performance on different system configurations.

OASB simulates realistic client scenarios using a selection of the most-commonly-used Oracle applications modules. Definitions of transactions that compose the benchmark load were obtained through collaboration with implementation consultants, and are representative of typical client workloads, with OLTP, batch, and self-service components.

The database used in the benchmark represents amounts of information that are typical of mid-market businesses whose annual revenues range from \$100 M to \$500 M. This database is provided with the benchmark kit and is common to all platforms on which the benchmark is available.

The user count measures the number of concurrent Oracle application users that the system can sustain while response times are kept under a pre-defined maximum value. User processes are defined by the type of transactions they execute, and each user maintains a minimum transaction-per-hour rate. Neither transaction rates nor the workload mix vary with increased system load or response times.

Why is Oracle Application Standard Benchmark important

Oracle applications handle many types of business needs, comprising of a complete suite of Enterprise Resource Planning (ERP) applications. These are usually implemented in conjunction with a re-engineering of business processes, and can be tailored to meet different requirements.

Oracle applications can be grouped into product families as follows:

- Financials
- Manufacturing
- Human resources
- Supply chain management
- Market management

The transactions that define OASB workloads were developed in close collaboration with database consultants. They are intended to represent a typical client workload with batch processing representing about 25% of the workload.

The OASB workload consists of the following components:

- Oracle Financial Applications
- Accounts Payable (AP)
- Accounts Receivable (AR)
- Fixed Assets (FA)
- General Ledger (GL)
- Oracle Supply Chain Management Applications
- Inventory (Inv)
- Order Entry (OE)
- Purchase Orders (PO)
- 18 OLTP transactions
- 7 batch jobs

Characteristics and constraints

OASB is a 3-tier benchmark that is intended to measure the performance of a realistic Oracle applications environment. The system under test comprises of application and database tiers, with simulated users driving the workload.

The middle tier handles the presentation and application logic for the user interface, connecting directly to the Oracle database. Like the SAP SD 3-tier benchmark, the workload is highly scalable across the application layer, so it is important to look at performance figures for systems that are a reasonable size for your environment.

A key influencer of performance in this benchmark is memory capacity and throughput in the application layer. A typical 4-way application server in this benchmark is configured with 32 GB of memory, significantly more than 8 GB of memory generally seen in four-processor application servers used for the more CPU intensive SAP SD 3-tier benchmark.

3.5.6 SPEC CPU2000

The focus of this benchmark is to compute intensive performance, which means that the main components that are tested in this benchmark are:

- CPU
- Memory architecture
- Compilers

SPEC CPU2000 measures system speed and throughput for single processor, symmetric multiprocessor, and cluster systems. A higher score in SPEC CPU2000 benchmark means better performance on the given workload.

SPEC CPU2000 is made of two sets of benchmarks, CINT2000 and CFP2000. They are measuring integer performance and float point performance and those are in turn divided in throughput metrics values and speed metrics values.

- CINT2000
  - It measures compute intensive integer performance.
  - The throughput metric is SPECint_rate_base2000, and it measures the number of tasks a computer can complete in a given amount of time.
  - The speed metric is SPECint_base2000, and it measures how fast a machine completes the running of the CINT2000 suite.
- CFP2000
  - It is used for compute intensive floating point performance.
  - The throughput metric, SPECfp_rate_base2000, measures the number of tasks a computer can complete in a given amount of time.
  - The speed metric, SPECfp_base2000, measures how fast a machine completes the running of the CFP2000 suite.

Run and reporting rules permit baseline and optimized peak results for the CINT2000 and CFP2000 suites. The metrics above states the required baseline results that must be a part of every SPEC CPU2000 benchmark.

For more details of the SPEC CPU2000 benchmark, visit the following Web site:

http://www.spec.org/cpu2000/

Why is SPEC CPU2000 standard benchmark important

The SPEC CPU2000 industry standard CPU benchmark suite provides a comparative measure of CPU intense performance across the widest range of hardware and different platforms.

SPEC CPU2000 is a highly portable benchmark. This makes it possible to run same code base compiled on a different platform and achieve the same relative result even if the platforms are different.

This benchmark is very good in cases where the CPU needs to be selected with specific characteristics. If a client wants high performance in floating point operations they would look for a CPU that has good scores in that metric.

Characteristics and constraints

SPEC CPU2000 is made for CPU benchmarking. It stresses the CPU and memory subsystem extensively. This benchmark does not stress other components in a server such as disk or networking.

SPEC provides benchmarking source code to be used with the benchmark. This code may note be altered.

3.5.7 SPECWeb99 and SEPCWeb99_SSL

Overview

SPECWeb99 measures the maximum number of simultaneous connections executing a given workload on a Web server, while still meeting specific throughput and error rate requirements. The connections are made and sustained at a specified maximum bit rate with a maximum segment size.

The SPECWeb99 is actually two benchmarks, SPECWeb99 and SPECWeb99_SSL. The major difference is that SPECweb99_SSL adds SSL support to the existing SPECweb99 workload. In SPECWeb99_SSL, the client systems generate the same mix of static and dynamic requests as SPECWeb99; however, the clients negotiate SSL connections with the server and all requests and responses are encrypted.

The performance result is expressed in the number of connections to the server for SPECWeb99 and the number of secure SSL connections to the server for SPECWeb99_SSL. According to SPEC, the results are not allowed to be compared. However for the same system, the differences between the SPECweb99 and SPECweb99_SSL results could be used to help evaluate the performance impact of SSL encryption on the server.

Both benchmarks will be replaced by a new web server benchmark that is under development called SPECWeb2004. It will consist of different workloads (both SSL and non-SSL), such as banking and e-commerce. The dynamic content is written in scripting languages to more closely model real-world deployments. The web server communicates with a lightweight back end to simulate an application/database server.

For more details of the SPECWeb99 and SPECWeb99_SSL benchmarks, refer to the following Web sites:

http://www.spec.org/web99/
http://www.spec.org/web99ssl/

Why is SPECWeb99 Standard Benchmark important

The SPECWeb99 is an industry standardized benchmark. Members of the SPEC organization—computer vendors, system integrators, universities, and many more—together agreed to the specification of this benchmark. The outcome of this is one standardized implementation and one standardized workload.

This is a benchmark that allows testing of different hardware platforms and Web server software to be tested in a way that allows an independent.

Characteristics and constraints

Neither SPECWeb99 or SPECWeb99_SSL takes into consideration the latency issues that may occur in a wide area network (WAN) or the Internet. This kind of behavior is very hard to simulate.

3.5.8 SPECjAppServer

The latest revision of SPECjAppServer is called SPECjAppServer2004 and it is a client/server benchmark that measures the performance of Java Enterprise Application Servers in a complete end-to-end Web application. SPECjAppServer2004 measures the scalability and performance of J2EE servers and containers that are based on the J2EE 1.3 specifications or later.

The latest revision is replacing the older SPECjAppServer2001 and SPECjAppServer2002 benchmarks. Some of the major differences between SPECjAppServer2004 and the older benchmarks is that they were based on J2EE 1.2 or earlier specifications.

Workload design is also one of the major changes with the latest revision of the benchmark. The workload that SPECjAppServer2004 uses emulates an automobile manufacturing company and its associated dealerships. The workload is provided as a kit with source code that is required to execute on the system if you wish to publish results. As a general rule SPEC does not allow modification of the code.

The SPECjAppServer2004 performance is measured by a metric called jAppServer Operations Per Second or JOPS. The metric is derived by adding the operations per second in the dealer domain of the workload to the work orders per second in the manufacturing domain.

Like other SPEC results, SPECjAppServer2004 results are reviewed by the SPEC benchmark committee (a technical peer review) and not audited independently like TPC. The SPEC Java subcommittee reviews the SPECjAppServer2004 results.

The result may not be compared with former revisions of the SPECjAppServer benchmark, and it may not be compared with TPC benchmarks either because they use different metrics among other things. The result may not be estimated or extrapolated either because of the complexity of the benchmark.

For more details of the SPECjAppServer2004 benchmark, visit the following Web site:

http://www.spec.org/jAppServer2004/

Characteristics and constraints

SPECjAppServer2004 primary objectives are to stress the J2EE application servers rather than the client tier or the database server tier. But other components are stressed since this is a solution based benchmark, like the database server.

This benchmark requires that the emulator (the function that is stressing the application server) is executed on a machine outside the system under test (SUT).

Following are the most significant influences on the performance of this benchmark:

- The hardware configuration
- The J2EE application server software
- The JVM software
- The database software
- JDBC drivers
- Network performance

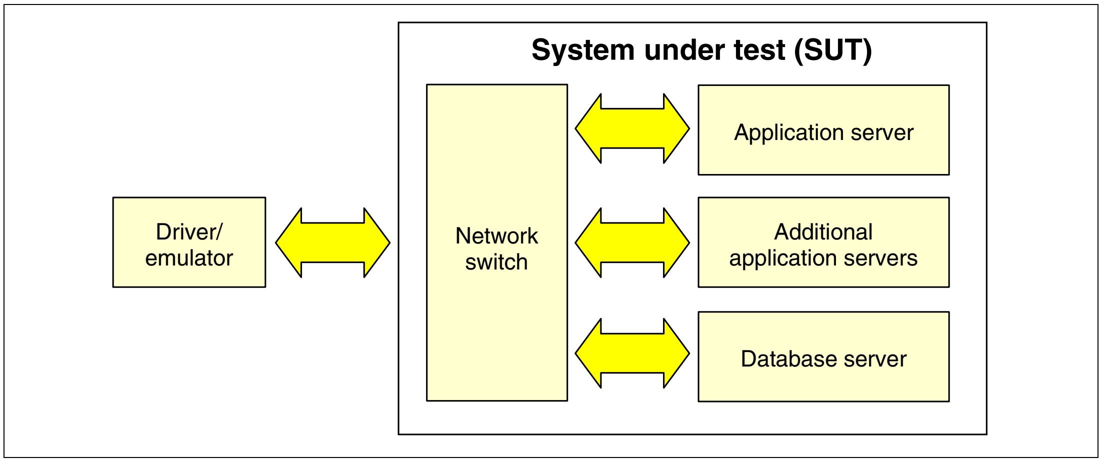

Figure 3-4   Architecture of the SPECjAppServer2004 benchmark

SPEC designed the SPECjAppServer2004 benchmark without explicit restrictions on scalability; however, it works well in both scale-out and scale-up configurations. How well it scales in any particular configuration is highly dependent on the underlaying hardware and software.

3.5.9 SPECjbb2000

SPECjbb200 is a benchmark that emulates a full 3-tier Java system that focuses on the middle tier on one server. This benchmark does not measure the network or database performance because that requires 2-tier or 3-tier configuration. SPECjbb2000 is totally self-contained and self-driving. SPECjbb2000 does not have support for clustered environments because it was not designed for it.

Measurement is done on the performance of CPUs, caches, memory, and the scalability of SMP platforms.

The metrics that SPECjbb2000 uses is ops/second. It is a throughput measurement representing the averaged throughput over a range of points during the benchmark process.

For more details of the SPECjbb2000 benchmark, visit the SPEC Web site:

http://www.spec.org/jbb2000/

Why is SPECjbb2000 standard benchmark important

Hardware and software vendors can use the SPECjbb2000 benchmark results to analyze their platforms scalability when running Java applications. Software vendors can also evaluate the efficiency of the JVMs, JITs, garbage collectors, and thread implementations of their applications.

Characteristics and constraints

This benchmark is rather modest in the requirement of hardware and it runs on one server.

SPEC provides the source code for the benchmark but requires that you run on the jar files they provide. If you recompile the code the result is invalid.

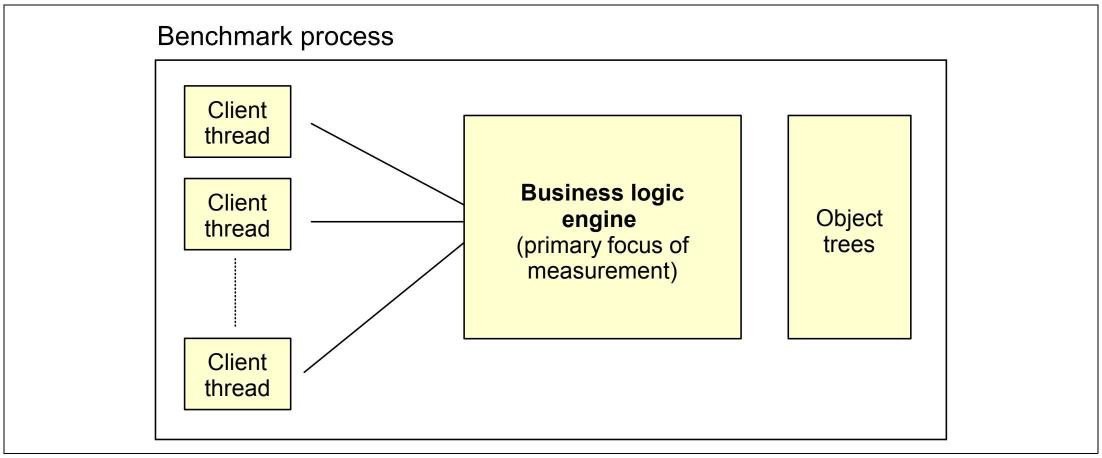

Figure 3-5 Schematics of the SPECjbb2000 benchmark architecture

4

Understanding benchmark results

You can use the results from industry standard benchmarks to help make informed decisions when assessing hardware and software platforms. Unfortunately, there is no magic formula for using these to evaluate system performance, as a standardized test cannot be expected to provide specific answers for a customized environment. The usefulness of any performance data is determined by how well you understand your own system and appreciate which pieces of benchmark data are relevant and which are not.

The following chapter includes information to help you understand a benchmark result, and to highlight some key considerations.

Following are the topics in this chapter:

4.1 Examining a benchmark result

Before examining industry standard benchmark results, there are three things that you should understand.

1. Understand your workload
2. Understand the workload that a benchmark is simulating
3. Understand how the benchmark relates to your workload

Without a reasonable understanding of these three key considerations, it is difficult to interpret a benchmark result in a meaningful way. There is also a fourth consideration, sometimes overlooked, which is to understand what is reasonable for your environment. Purchasing a new vehicle is a good analogy to illustrate these elements:

> If you plan to buy a new vehicle, looking at engine specifications or how fast it can accelerate from 0 to 60mph is not useful until you know its purpose. These may be important considerations if you are buying a motorcycle to race, but if your requirement is to transport ten children to school, factors such as passenger capacity, are likely to be more important. Taking this analogy one step further, you may decide that although it easily performs the task, a school bus is not a *reasonable* option, so you limit your search to vans.

Similarly, before looking at benchmark results, it is necessary to understand what purpose a server is used for, and what are the characteristics of the application workload. Once you understand this then you can examine the benchmarks that best represents the workload. 3.3, “Selecting the relevant benchmark” on page 23, can assist in defining the workload and determining the most appropriate benchmarks to consider.

4.2 Finding information

The industry standard benchmark owners publish results on their websites. Refer to 3.5, “Benchmark descriptions” on page 27 for the Web site of each benchmark. Generally the benchmark results are listed in order, beginning with the highest performance result to the lowest. Additionally, the results can usually be searched or sorted in different ways.

As an example, Figure 4-1 on page 41 shows the sorting options for the TPC-C and TPC-H benchmark results on the TPC website. The top ten results can be shown, sorted by performance or by price/performance, and if preferred, limited to clustered systems or non-clustered systems. There is also the option to view all the results, sorted by hardware or database vendor.

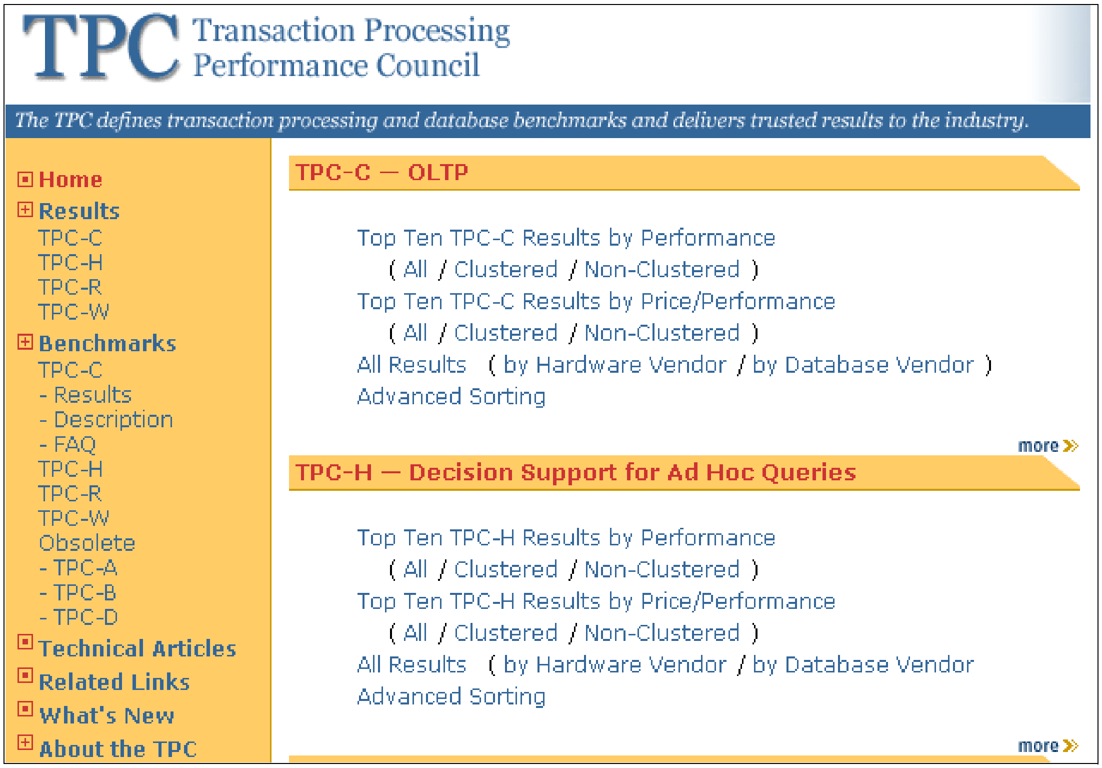

Figure 4-1 Benchmark results viewing options

When looking at these results it is important to understand a little about the benchmark itself. The highest performance figure is not necessarily the most appropriate result to look at. Selecting the top ten list of TPC-C benchmarks, most likely displays some very large 64-way servers.

This is where you consider what is reasonable in your environment. A 64-way server may be reasonable for your environment, and if so, these may be relevant benchmarks to examine. If a server with 4 or 8 processors is more appropriate, then those are the results you should find and examine. Do this by listing all results sorted by hardware vendor, or by using the advanced sorting option.

As a part of understanding your workload, it is important to appreciate the restrictions and limitations within your environment. Existing assets may have tied you to a particular piece of hardware or software, even to a particular version or release. Users and administrator skills may determine whether a major change of platform is possible. These factors may limit which configurations are viable, and help you to select which benchmark results to examine. For example if you have only Microsoft Windows trained staff, and standardized on Intel's EM64T processors, then this may limit your search to a small subset of the available benchmarks.

4.3 Disclosure reports

Virtually all of the industry standard benchmarks require a full disclosure report to be submitted with a result. The level of detail required varies between benchmarks, but generally provides an enormous amount of information regarding the system configuration. Full disclosure reports are intended to describe all components of the benchmarked system, including the hardware, software, and every parameter setting required to reproduce the result.

Do not take the reported performance metrics of a benchmark on face value alone. The disclosure reports provide basic configuration information that is vital for apples-to-apples comparisons. That is, information that allows you to ensure you compare equivalent systems, or that you understand the differences in the systems compared. Some of the basic configuration information provided may be:

- Number, frequency, and cache of processors
- Amount and speed of memory
- Number and speed of disks
- Network interfaces

In addition there is much more information that may be very useful, sometimes more useful than the performance metrics themselves. Depending on the benchmark, there are many other options and settings that are documented in the disclosure.

- Operating system
- Database
- Single system or clustered system
- Number of servers and configuration used in the cluster
- Compiler
- Compiler options used
- Disk configuration
- Tuning parameters

4.4 Workload characteristics affect system performance

Different workloads stress servers in different ways, tending to bottleneck different subsystems. A very good illustration of this appears when examining an xSeries 445, configured for different benchmarks. Table 4-1 shows the memory and disk configurations for an eight processor xSeries 445 used for seven different benchmarks.

Table 4-1 xSeries 445 configured for different benchmarks (8-way categories)

Benchmark	Processors	Memory (GB)	Number of Disks
TPC-C	8	64	409
TPC-H	8	32	138
SAP (SD 2-Tier)	8	16	15
Oracle (OASB 11.5.6)	8	32	42
Peoplesoft GL 8.4	8	8	92
SPEC CPU2000	8	4	2
SPECjbb2000	8	32	1

The different server configurations reflect the different workloads of the these benchmarks. The workload that the benchmark generates, causes the server to bottleneck in a particular subsystem. To alleviate the problem, resources are added to that subsystem, until the bottleneck is moved to another subsystem. The flowchart in Figure 4-2 on page 43 shows a very simplistic way to look at how a system is configured for a benchmark. For most benchmarks the process is completed when the bottleneck of the system is the CPU operating at close to 100% utilization.

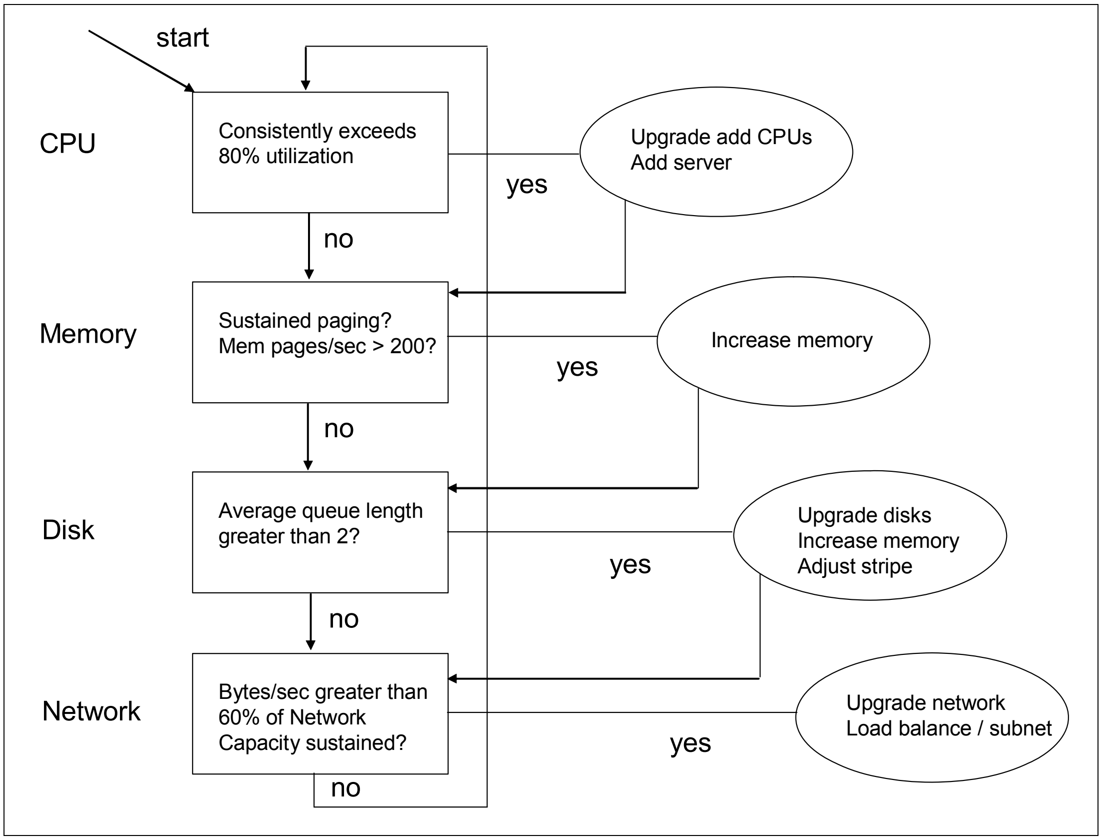

Figure 4-2 Bottleneck detection flowchart

Following this logic, the server configurations listed in Table 4-1 on page 42 indicate roughly how each benchmark stresses the CPU, memory, and I/O subsystems. The TPC-C benchmark heavily stresses all three subsystems, Oracle OASB performance is heavily influenced by memory, and Peoplesoft General Ledger benchmark has a very high I/O requirement.

The configurations in Table 4-1 on page 42 also highlights the component focused nature of the SPEC benchmarks. For example the 8-way xSeries 445 used for the SPEC CPU2000 benchmark, required only 4 GB of memory and two disks. Clearly the workload isolates the CPU with very little dependency on other subsystems. This means that the benchmark may be very good for comparing raw CPU performance, but it provides limited information regarding the performance of the entire system. The CPUs in a system may be very fast, but performance remains poor if the memory or I/O subsystems cannot supply data to them quickly enough.

Performance theory: When benchmarking a server, the ultimate goal is to make the fastest component in the system operate at 100% utilization. In modern servers, the CPU is the fastest subsystem. Each of the configurations listed in Table 4-1 on page 42 saturate the CPUs for that particular benchmark. For the TPC-C workload this required the 8-way x445 to be configured with 64 GB of memory, and in excess of 400 disks. The SAP SD 2-Tier benchmark, a more CPU intensive workload, required 16 GB memory and 15 disks.

4.5 Types of benchmarks

It is important to understand what a benchmark is testing, and what are its limitations. There are two main categories of benchmark: component level and system level.

4.5.1 Component level benchmarks

Component level benchmarks (also called synthetic benchmarks) attempt to isolate one specific component, such as the processor, and to comprehensively test that component in order to evaluate its performance characteristics. These benchmarks are very good for singling out individual components for direct comparison. For the purpose of this paper, this category also includes benchmarks that isolate a single task or function, such as SPECjbb2000. Although this benchmark does model a 3 tier system, within a single server it is essentially a benchmark of CPU, memory, and Java Virtual Machine (JVM) performance, having no network or I/O requirement.

These types of benchmarks are sometimes criticized and discounted as being too far from real-world applications. This is partially true in that you must take caution when using the results, as they do not take into account other factors that affect performance. For example, SPECjbb2000 does not have any I/O requirement, and the database is replaced by a binary tree of objects. In a real-world environment these factors exist and may heavily influence the performance of a system.

However, for applications that generate workloads that heavily use the component tested, a component level benchmark provides valuable information. Component benchmarks tend to require more knowledge and analysis in order to come to meaningful conclusions, because your own system performance is likely to be impacted by system elements not tested.

The SPEC benchmarks tend to fall into this category of component level benchmarks. There are exceptions to this, such as SPECjAppServer, and in the future, more SPEC benchmarks are likely to test entire systems with workloads that are representative of real-world applications.

4.5.2 System-level benchmarks

System-level benchmarks generally emulate a simplified real-world computing environment. They are easier to understand at a high level, as they are described in terms of an application workload, and performance takes into account multiple aspects of a system. These are very useful benchmarks to consider, but can be very inaccurate if the workload characteristics do not resemble the workload for which you are testing.

The TPC benchmarks fall into this category of system level benchmarks, along with most of the application specific or ISV benchmarks.

4.6 Benchmark result considerations

When you examine a benchmark result, there are a number of important aspects that may not be apparent from looking at the performance metrics alone. This section discusses some of those considerations.

4.6.1 Know what you are comparing

In order to evaluate and compare system performance in a meaningful way, compare equivalent systems. We mentioned this several times throughout this paper, but it is an extremely important point. For example, one vendor may claim superior performance of their server, Model A, versus a competitor's server, Model B. They produce performance metrics from an industry standard benchmark to support the claim. Although the servers were configured in a similar manner, inspection of the full disclosure report shows that the Model A server used 3.2 GHz processors with 2 MB L2 cache, in comparison to 3.2 GHz processors with a 1 MB L2 cache for Model B server. This difference alone for many benchmarks may be more than enough to account for the superior result, so the comparison is not valid.

4.6.2 Maximum performance versus price/performance

Systems are not always benchmarked to produce a maximum performance figure. A number of industry standard benchmarks require the total cost of the system to be published, so vendors sometimes also benchmark systems configured with less memory, less disk, or less expensive applications, in order to produce a best price/performance figure. Unfortunately the disclosure reports do not state that a particular benchmark performed for price/performance, rather than maximum performance. Generally this is seen by obvious differences in performance metrics and price/performance metrics, but you can miss it if you look at performance figures alone. A good practice is to always look at the disclosure and make sure that configurations are similar.

4.6.3 Clustered system benchmarks

It is important to understand what a benchmark result produced with a clustered system signifies for a particular benchmark. The workloads of many benchmarks are highly partitionable, allowing the processing to be spread over a number of clustered servers.

This is very useful in demonstrating the ability of a database or application to scale-out. Scale out environments have become very popular as they allow capacity growth by adding low cost servers.

However, this can lead to a deep pocket benchmark. This means that by adding more and more servers, the performance figure is pushed higher and higher. There is a limit, but it may be an extraordinarily large configuration, that is totally unrealistic in a normal environment. This is not to say that the information provided by the benchmark has no value, but it does mean that care must be taken to look at benchmark results where the system represents a realistic configuration for your workload.

The SAP Sales and Distribution (SD) benchmark provides a good example to illustrate this. There are two forms of the benchmark, 2-tier and 3-tier. In the 2-tier version of the benchmark the SAP application and database are installed on the same server. The application is very CPU intensive, and during the benchmark consumes approximately 80% of the CPUs. The database consumes only about 10% of the CPU resources.

The SAP SD 3-tier benchmark allows the application layer to be separated out to additional servers, which is a common architecture in real-world environments. This is shown in Table 4-3 on page 46.

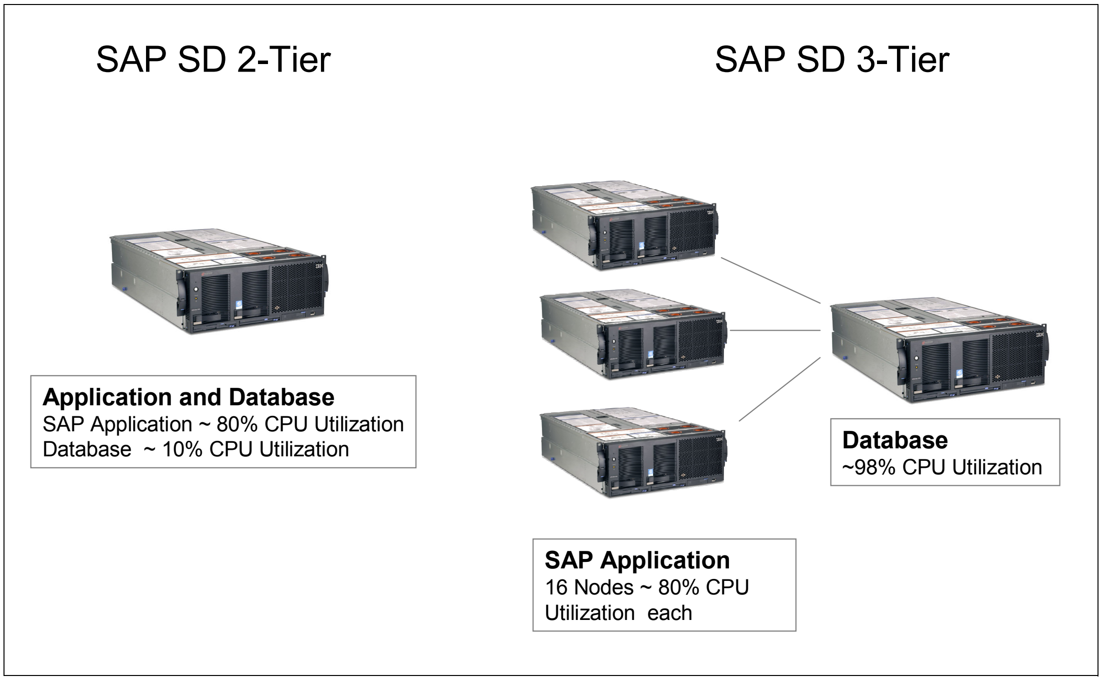

Figure 4-3 SAP SD 2-tier and 3-tier configurations

Performance in the SAP benchmarks is related to the number of users the system can support with response times of less than two seconds. This is increased by adding more and more application servers until the database server becomes saturated. That is, the CPU utilization reaches close to 100%. However, databases tend to scale extremely well, so you can increase the size of the database server to allow even more application servers to be added. There is a limit to this, but it is a very large and unrealistic configuration. Results for this benchmark are posted from configurations consisting of a 64-way database server, and 160 application servers.

This does not mean that the SAP SD 3-tier benchmark has no relevance. The scale-out nature of the application layer is common in commercial architectures, and the benchmark highlights this. However, when looking at the benchmark results, examine results where the configuration is reasonable for your environment. Many realistic configurations are submitted for this benchmark, for example, the application layer contained in one BladeCenter chassis.

4.7 Benchmark results as part of an evaluation process

As discussed throughout this Redpaper, the performance results achieved by vendors using standardized benchmarks are unlikely to correspond exactly to your environment. If this is the case, then what is the point? How are they useful to assist evaluating hardware and software platforms?

The industry standard benchmarks, such as those defined by TPC and SPEC, are developed as collaborative projects amongst a wide variety of vendors, so the tests are not likely to be biased towards any particular vendor or configuration. The benchmark descriptions are usually available to anyone so that it is possible to understand which parts of a benchmark may apply to your environment. The rules that accompany these standardized tests also require a fair degree of disclosure making it possible to learn what features are required to achieve the high levels of performance. These standardized benchmarks also provide a level playing field where everyone agreed to compete fairly. Even if the playing field is not precisely the same as your own environment, the results may provide insights into inherent strengths and weaknesses of different vendors' offerings.

The task in understanding a benchmark result, is to extract the data from these sources and determine what they mean for your environment. Extract all the relevant data from the standardized tests and get a full disclosure of the configurations tested to be sure that the results quoted are applicable to your configuration. Finally, be sure to evaluate all data as weighed by its relevance to your own workload. A dramatic benchmark result may not help you much if it does not stress any of the system that you care about.

5

Client-related benchmarks

Why would any client want to perform their own benchmarking? The two typical reasons are to compare systems they are considering purchasing or to better understand the performance of their own system, perhaps before or after a system upgrade.

In this chapter we discuss the best practices that a client can do to make a benchmark as easy as possible. We also examine some of the most common traps and pitfalls associated with performing a benchmark. However, this chapter does not explain how to perform your own benchmarking.

The topics covered are:

5.1 Industry benchmarks versus client benchmarks

A benchmark is a standardized problem or test used to measure system performance, but what are the major differences in requirements between an industry standard benchmark and a client benchmark? An easy comparison appears in the following table.

Table 5-1 Differences in requirements between a industry benchmark and a client benchmark

Values to evaluate	Requirements for industry benchmarking	Requirements for client benchmarking
Soft values		
Involvement of staff	High	Low to moderate
Technical skill	Very High	High
Tuning skill	Extremely high	High
Time consumption	Very high	Moderate
Cost for benchmarking	Moderate to extremely high	Moderate to high
Publication of result (FDR)	Extremely high	Low
Benchmark reruns	Moderate to very high	Low to moderate
Understand the character of the system as a unit (hardware and software)	Very high	Very high
Hard values		
Server requirements	High to very high	Low to high
Disk requirements	Low to extremely high	Low to moderate
Software requirements	Moderate to high	Moderate
Workload application requirements	Moderate to extremely high	Moderate

There are a few industry standard benchmarks clients can reasonably expect to be able to perform on their own. One example is SPEC CPU2000. As we described in 3.5.6, “SPEC CPU2000” on page 34, this is a benchmark that considers only two subsystems in a server, CPU and memory.

As we previously mentioned, performing industry standard benchmarks is not easily done. There are some industry benchmarks that clients can execute, but they tend to stress specific subsystems in a server and not the entire system or infrastructure as a whole unit. We strongly discourage clients from performing industry benchmarks on their own. As we can conclude from Table 5-1, the requirements are too high for the most industry standard benchmarks.

5.2 Alternatives to client benchmarking

There are a number of ways a client can gather performance information about systems being compared before making a selection. We describe and show pictorially in order of increasing complexity in Figure 5-1 on page 51.

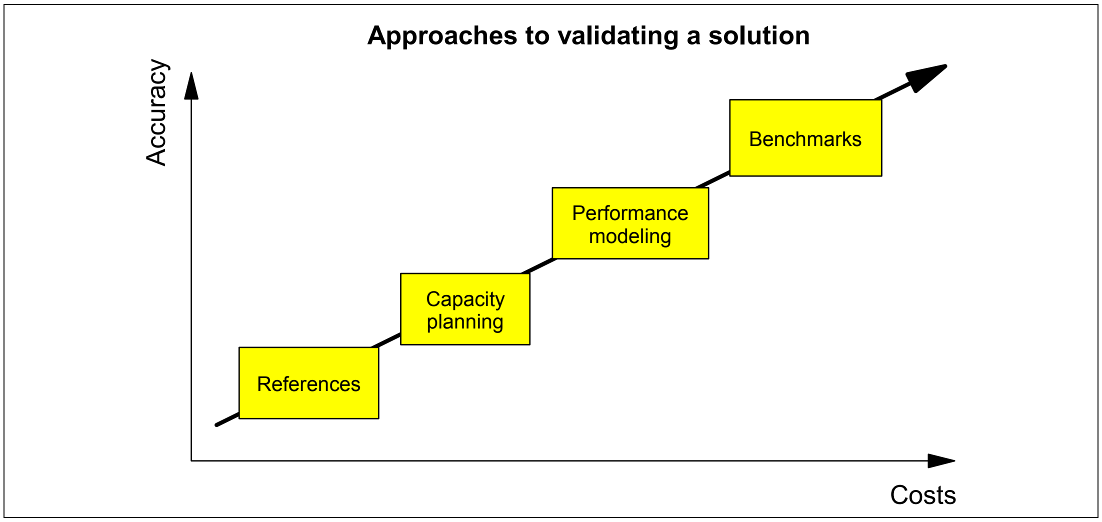

Figure 5-1 Different methods of validating a solution

- References

  With a reference visit, a client can see another client's environment that is very close to their own requirements. This is the easiest way for a client to get input but also the one with the least accuracy. The referencing setup may be close but how do the workloads compare to each other?

- Capacity planing

  The use of capacity planning tools, such as IBM Director's Capacity Manager, enables clients to extrapolate future component requirements on their existing servers. The accuracy of this method is not good if the client wants to replace the existing servers and the further out you extrapolate the less accurate it is. In the absence of any other analysis it gives rudimentary indicators as to future requirements.

- Performance modeling

  With this method, the workload is divided into separate components such as RAM, disk, and CPU. This is a difficult exercise and requires a significant understanding of the applications being used to test. However, if done properly, accuracy is high.

- Benchmark the actual application

  Benchmarking is the best way to test the setup because a workload is placed on the system that gives us an understanding of how much client load the system can take before bottlenecks start appearing and levelling off performance. It also gives us a good understanding on how we can balance the system, and how we can grow with it by adding different components. The accuracy is highest with benchmarks but the cost is also the greatest.

  An option that may be available to clients is to request that IBM Advanced Technical Support undertake custom benchmarks on their behalf. For more information, IBM employees can access the following Web site:

  `http://w3.ibm.com/support/americas/bench/`

5.3 Why perform your own benchmarking

Following are some of the reasons why you would perform your own benchmarking:

- You do not use any of the applications used in industry standard benchmarks, so the results do not relate to your environment. Not all clients use major applications that vendors produce industry standard benchmark results for such as SAP, BAAN, Microsoft SQL Server, Oracle, Peoplesoft, or Lotus Domino.

  Many clients developed their own custom applications that are the core of their business infrastructure. These in-house developed applications suit them better than using off-the-shelf applications. These applications are typically only used by the one company and in such cases, it is hard to compare their performance characteristics with the industry standard benchmarks.
- The hardware configurations you are considering do not match the results of industry benchmarks, and you feel that the only way to make a valid comparison is by performing your own analysis. You may, for example, have a different CPU and memory setup for which no results are published.
- You want to find optimal price/performance figures for your specific configurations or configuration changes. For example, you may want to know if putting 2GB of RAM in a system instead of 1 GB gives you a sufficient gain in performance to be worth this investment.

5.4 What to expect if you must do your own benchmarking

Performing a benchmark is an not an easy task. It is a large and time consuming process.

When you decide to perform your own benchmarks because there is no other way to get the information that is needed, there are a couple of things that are very important to consider:

- What is the goals of performing the benchmark?
- What do you expect to learn from it?
- How much time are you ready to put aside for this?
- How much money are you prepared to invest?

A client-related benchmark should have, as any other project, a goal, a time line, and a budget. Without those cornerstones, the project may fail and end up with a lager amount of money spent, the goals not achieved, you learned less then you expected, or even worse, all of the above.

5.5 Traps and pitfalls of doing your own benchmarking

Based on the experience of the engineers in the IBM Performance Lab, following are some mistakes you may make when attempting to perform your own benchmarks:

- Using a desktop benchmarking tool when validating servers

  A very common mistake is to use benchmarking tools that are intended for desktop and workstation environments. Those kind of tools are not made for simulating a workload on a server. They are normally used to test a subsystem in a workstation like the graphic subsystem or the CPU subsystem. This usually results in a poor result that is not indicative of the server's true performance in a multi-user server environment.

The classic example is using XCOPY to copy a large file from a client to the server as a benchmark for server performance. Servers are meant for a large number of users typically effectively causing a random workload against the server.

- Using the wrong benchmark for your workload

  If your server or intended workload is running a decision support-type application, then performing an OLTP benchmark (or more likely, reviewing a TPC-C OLTP benchmark result) does not provide any meaningful data.

- Misunderstanding the terms transaction and user

  Two very common questions that clients ask are “How many transaction can this server handle?” and “How many users can this server support?”. This is a big source of confusion and errors.

  For example, if the clients ask about how many tpmC transactions a server can handle then the question is valid. tpmC is a well defined performance metric that TPC defined for the TPC-C benchmark, and it is expressed in transactions-per-minute-C or tpmC.

  If the clients ask in terms of how many transactions the server can handle in their specific environment, then the question is no longer valid. The clients must ask themselves what a transaction is for them. This is not very easy to state and must be handled with care.

  The same goes for the term *user*. If the client asks how many users an SAP server can support based on a standard SAP application benchmark then it is a valid question. The term user is well defined by SAP as user interaction repeated n times for a minimum of 150 second duration.

  If the client wants to know how many users their own system can support then it is not a valid question. The client must define their own user definitions based on the character of how their employees work against the system, They should define how a light user interacts with the system versus how a heavy user interacts with the system. Based on those definitions the client can perform their benchmarks, and then state how many users a system can support.

  So at the end it is a misunderstanding of definitions. There are no general definitions of the terms transaction and user. They must be defined if they are to be used as a metric that a client wants to use.

- Overestimated expectations of adding new and faster components to a system

  Very often, when a client adds new components to a server the expectation is that it will go much faster. For example, a client replaces hard drives in a server that runs at 10 000 RPM with ones that run at 15 000 RPM and expects that the disk subsystem will give them 50% more drive throughput because the new drives are 50% faster then the old ones. This is usually not the case.

  It is very important for a client to understand what factors can affect the performance of, in this example, the disk subsystem. These factors include:

  – RAID strategy
  – Number of drives
  – Active data set size (effectively, how far the drive heads have to physically move)
  – Disk stripe size
  – Drive speed
  – Disk cache write-back/write-through
  – RAID adapter cache size

  As we can see we have only affected one of those areas, the drive speed. So adding just faster drives does not necessarily increase the performance according to expectations.

  If the client wanted a 50% increase in drive throughput they must double the number of drives rather than buying faster drives.

- Faster or more CPUs always raise server performance

  A general misconception is that if a faster CPU is added to the system or if a server is turned into a SMP system with multiple CPUs the server will run faster. The server might run faster but usually not as fast as the client expects it to run.

  For many years now the performance development for CPU doubled every 18 months approximately (also known as Moores law). During this time the other components that are part of a server, like disk and memory, have not increased in speed at the same rate.

  So if the CPU subsystem is getting faster and faster and the other subsystems are not developing in the same rate, the only way to keep the server in balance is to add more disk and memory so they can keep up with the CPU subsystem.

  Figure 5-2 shows what happens when CPUs are added to a database server (OLTP-type application), and how this impacts on relative performance with and without also adding memory at the same time.

  If we add a second CPU to a server system and the memory is doubled we get a relative performance improvement of 1.74.

  If we add two more CPUs, there is a difference in relative performance improvement depending on whether we also add memory. Just adding two CPUs (from two-way to four-way) produces a relative improvement of 1.4; however, it also doubles the memory improvement performance by 1.61.

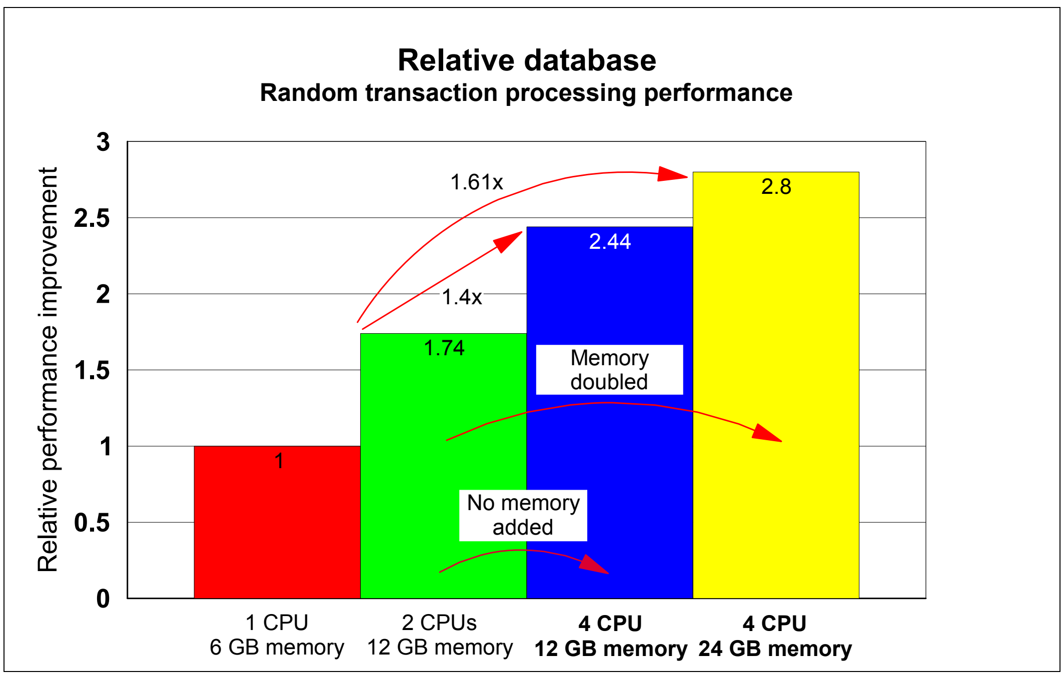

Figure 5-2 Relative performance scaling from adding 3.0 GHz processors

  In general, the largest system performance improvements are achieved when you add memory and disk drives. Rarely is the CPU the bottleneck in a sever. So upgrading the CPU simply leads to the system running with lower CPU utilization, while other bottlenecked components, such as disk and memory, become even more of a bottleneck.

- One hardware configuration fits across all environments and workload

  Clients sometime expect that the same server configuration will give the best performance and scalability for all types of workloads. This is not true. One size does not fit all solutions and workloads.

When maximizing system performance, you are trying to optimize every link to handle the workload as good as its surrounding subsystems. A system's performance is at its best when all subsystems are in balance and all of them are at their maximum capacity. We discussed this in Chapter 4, "Understanding benchmark results" on page 39, but it is worth repeating the discussion.

You can see a very good illustration of this can when examining the different configurations of an xSeries 445, for different benchmarks. Table 5-2 shows the memory and disk configurations for an eight processor xSeries 445 used for seven different benchmarks. You can see the effects of the different workloads by looking at the server configurations for these benchmarks.

Table 5-2 xSeries 445 configured for different benchmarks (8-way categories)

Benchmark	Processors	Memory (GB)	Number of Disks
TPC-C	8	64	409
TPC-H	8	32	138
SAP (SD 2-Tier)	8	16	15
Oracle (OASB 11.5.6)	8	32	42
Peoplesoft GL 8.4	8	8	92
SPEC CPU2000	8	4	2
SPECjbb2000	8	32	1

If we make a configuration for a TPC-C benchmark and one for TPC-H and the goal is to get a balanced system that has all subsystems running at full capacity, then the configurations will look different—the TPC-C configuration will need four times as many disks and twice the RAM as the TPC-H configuration for the same processors.

5.6 Summary

Following is a summary of this chapter:

- Ask yourself if there is another way besides benchmarking to get the data that you need in order to make decisions on new servers or new components. Benchmarking can be expensive in terms of time, effort, and money.
- If there is no other way, do not take benchmarking on your own lightly. It is a very big process, and it requires a lot of technical and performance tuning skills.
- Treat the benchmark process as a project with a goal, a time line, and a budget. By doing so you might save time and money and achieve actual results that you can use.
- If you want to perform a benchmark, you must understand the character of your total environment as well as all of the subcomponents.
- You must be able to put the correct workload on the systems and applications. If you do not understand your workload and cannot define it, get help from experts.
- Do not expect miracles in performance by adding more hardware to just one subsystem. Performance tuning is an advanced topic, and you must understand the interaction between different subsystems when upgrading subsystems.

Abbreviations and acronyms

CAD	computer aided design
CMT	Center for Microsoft Technologies
DSS	decision support systems
ERP	enterprise resource planning
EXA	Enterprise X-Architecture
GB	Gigabytes
GPC	Graphics Performance Characterization
HAL	hardware abstraction layer
HCT	hardware compatibility test
HPC	high performance computing
HPG	High-Performance Group
ISV	independent software vendor
JDBC	Java Database Connectivity
JOPS	jAppServer operations per second
JVM	Java Virtual Machine
LTC	Linux Technology Center
MAPI	messaging application programming interface
MB	megabytes
MIPS	millions of instructions per second
OASB	Oracle Applications Standard Benchmark)
OLTP	online transaction processing
OSG	Open Systems Group
PAE	Physical Address Extension
PSPP	Platform Sizing and Performance Program
RAID	redundant array of independent disks
RAM	random access memory
RDBMS	relational database management system
RPM	revolutions per minute
RSA	Remote Supervisor Adapter
RTE	Remote Terminal Emulator
SCT	system compatibility test
SIPS	service interactions per second
SMP	symmetric multiprocessing
SOAP	Simple Object Access Protocol
SPEC	Standard Performance Evaluation Corporation
SQL	structured query language
SSL	Secure Sockets Layer
SUT	system under test
TPC	Transaction Processing Performance Council
UDDI	Universal Description, Discovery and Integration
WAN	wide area network
WSDL	Web Services Description Language
XML	Extensible Markup Language

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 60. Note that some of the documents referenced here may be available in softcopy only.

- *Tuning IBM eServer xSeries Servers for Performance*, SG24-5287-03
- *Help Me Find My IBM eServer xSeries Performance Problem!*, REDP-3938-00
- *Tuning Windows Server 2003 on IBM eServer xSeries Servers*, REDP-3943-00
- *Tuning Red Hat Enterprise Linux on xSeries Servers,* REDP-3861-00
- *Tuning SUSE LINUX Enterprise Server on IBM eServer xSeries Servers,* REDP-3862-00

Online resources

These Web sites are also relevant as further information sources:

- ServerProven for xSeries

  http://www.pc.ibm.com/us/compat/indexsp.html
- Poughkeepsie pSeries & xSeries Benchmark Centers (IBM employee access only)

  http://benchmarks.pbm.ihost.com
- EMEA Product and Solutions Support Centre (IBM employee access only)

  http://w3.ibm.com/support/pssc
- Americas ATS Benchmark Centers

  http://w3.ibm.com/support/americas/bench
- Transaction Processing Performance Council (TPC) Web site

  http://www.tpc.org
- TPC benchmarks:

  TPC-C: http://www.tpc.org/tpcc
  TPC-H: http://www.tpc.org/tpch
- Standard Performance Evaluation Corporation (SPEC) Web site

  http://www.spec.org
- SPEC benchmarks:

  SPECweb99: http://www.spec.org/web99
  SPECweb99ssl: http://www.spec.org/web99ssl
  SPECjbb2000: http://www.spec.org/jbb2000
  SPECcpu2000: http://www.spec.org/cpu2000
  SPECjAppServer2004: http://www.spec.org/jAppServer2004
- Oracle Applications Standard Benchmark

http://www.oracle.com/apps_benchmark

- SAP Standard Application Benchmarks

  http://www.sap.com/benchmark

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Understanding IBM @server xSeries Benchmarks

Redpaper

What industry benchmarks are used in the xSeries marketplace

How to interpret the results of the benchmarks

How the benchmarks relate to client configurations

Many models of the IBM @server xSeries family maintained a leadership position for benchmark results for several years. These benchmarks help clients position xSeries servers in the marketplace, but they also offer other advantages to clients including driving the industry forward as a whole by improving the performance of applications, drivers, operating systems, and firmware.

There is a common misconception that industry benchmark results are irrelevant because they do not reflect the reality of client configurations and the performance and transaction throughput that is actually possible in the "real world". This Redpaper shows that benchmarks are useful and relevant to clients and that benchmark results are useful when attempting to understand how one solution offering performs over another.

The purpose of this Redpaper is to explain what benchmarks are and how to interpret the benchmark results so you can understand how they relate to their own server plans. The major industry benchmarks from the Transaction Processing Performance Council (TPC) and the Standard Performance Evaluaction Corporation (SPEC) are described, explaining how they relate to specific client application types. This paper is for clients, IBM Business Partners, and IBM employees who want to understand benchmarks on xSeries servers.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks