

Building a Scalable e-Business Platform

*A Performance Case Study for
WebSphere™ Application Server
and
DB2® Universal Database™*

October 2001



Introduction

The demise of dot com mania has made everyone keenly aware that deploying e-commerce technology doesn't create business out of thin air, it merely provides a platform for conducting business more efficiently. It is clear e-commerce has fundamentally changed the speed at which business is conducted – and established a near zero tolerance for poorly performing e-commerce applications. To effectively compete, companies must be able to interact with an ever changing set of suppliers, customers, and colleagues – and do it well. This demands an infrastructure that offers an “enterprise-class” level of performance, including high levels of availability and scalability. Increasingly, this infrastructure is centered around dynamic and data-driven web based applications.

In today's fluctuating economy, any strategic investment must be given careful consideration. Building an information technology infrastructure that is only focused on serving the needs of today, could spell disaster for tomorrow. Companies deploying e-commerce must insure that the building blocks of that infrastructure will also expand to meet their future challenges, maximizing the return on their investment. In short, they want a cost-effective platform that will meet growing demands. Today's business environment carries enough risk and uncertainty without building an e-business infrastructure on unknowns. It is important to have an infrastructure that will grow in a predictable manner to serve the dynamic business demands

This paper documents the highlights of a joint effort by IBM, Quest Software, and Intel Corporation to demonstrate how IBM's e-business software, WebSphere and DB2, running on IBM's Intel-based servers, provides such an infrastructure. A set of robust benchmark tests, designed to emulate an on-line brokerage firm, were conducted at IBM's xSeries Teraplex Integration Center to measure the scalability characteristics of representative configurations as they grow to service more users and more transactions. Quest Benchmark Factory was used to simulate volumes of customers accessing trading services via web browsers. Test results demonstrated convincingly that WebSphere and DB2 running on IBM Intel-based servers deliver a scalable, cost effective solution with the predictability that is critical to manage growing e-business systems.

Project Objectives

IBM, Quest Software and Intel Corporation joined forces to conduct a performance study of the scalability characteristics required of an e-business infrastructure. The objectives of the study were to:

- establish proof points for the predictable scalability of IBM's WebSphere Application Server and DB2 UDB for e-business solutions,
- demonstrate the viability of IBM's Intel-based xSeries servers for scalable e-business solutions and,
- demonstrate the robustness, flexibility and versatility of Quest Software's Benchmark Factory as a tool for measuring and capturing scalability characteristics to be used in capacity planning of web-based transaction environments.

Selecting the Ingredients for Success

A project of this scope required the building of a focused test environment on a system infrastructure based on the integration of industry-leading technology, a robust and versatile measurement capability, a team of people skilled in planning, developing and executing a complex performance study and a large scale testing facility. The following section describes the technology infrastructure, test environment and the skills and facilities that made up the final project.

IBM's WebSphere Application Server and DB2 Universal Database

IBM's WebSphere Application Server and **DB2 Universal Database** software technologies were chosen as the web and database building blocks respectively, and IBM's xSeries servers, based on Intel processor technology provided the processing hardware for the infrastructure. **WebSphere Application Server**, Advanced Edition, provides a strong, cross-platform, Java technology-based Web application platform capable of supporting medium-to high-level transactional environments used in conjunction with dynamic Web page serving. Known for its robust functionality and high reliability, **DB2 Universal Database** offers the data integration capabilities that enable e-businesses.

IBM's xSeries Servers Based on Intel Technology

IBM's xSeries servers incorporate mainframe-inspired technologies and intelligent management tools for real hands-off reliability. Part of the IBM hardware portfolio for e-business, xSeries systems are a high-performance line of Intel-based servers spanning the solution spectrum. The xSeries model x370 used in this study, part of the rack mount category of models and can be configured with up to eight high-performance Intel® Pentium® III Xeon processors, gives you the ability to process large amounts of data while getting extreme world class system performance. As the world's leading supplier of processors designed for servers, Intel has developed a powerful suite of server building blocks and tools that enable original equipment manufacturers (OEM), local system integrators and resellers to build world-class servers.

WebSphere Performance Benchmark Sample, a.k.a. "Trade2", e-business Application Simulator

Derived from experiences with many customer environments, **Trade2** is a collection of Java classes, Java Servlets, Java Server Pages and Enterprise Java Beans integrated into a single application, designed to simulate an online brokerage firm. Originally built by the WebSphere Performance team as an IBM internal tool for testing the performance of the WebSphere Application Server, Trade2 is now available to IBM's customers to download from the Internet. More information on this benchmark can be found in the last section of the paper.

Quest Software Benchmark Factory Load Testing Software

Benchmark Factory was selected as the workload simulator primarily for its scalability and ease of use. Benchmark Factory is load testing and capacity planning software that scales throughput testing to virtually unlimited users. Benchmark Factory prevents system downtime and poor performance by determining the critical limits of any application environment before the environment is live.

Team of Experts and IBM xSeries Teraplex Integration Center

With years of success testing business intelligence solutions, integrating IBM system products with non-IBM components, **IBM's xSeries Teraplex** Center in Raleigh, North Carolina, was ideally qualified to provide an integrated and flexible hardware and software environment, as well as a team of specialists with the project management and technical skills necessary to support a project of this scope and complexity. Performance experts from IBM's DB2 e-business Performance team in Toronto were recruited to develop a comprehensive test plan, create and execute complex test scenarios, take measurements and analyze results. Quest Software provided on-site installation and customization support for Benchmark Factory.

More information about each of the technologies, products and the Teraplex test facilities can be found in the Solution Details section of this document.

Benchmark Configuration

To effectively demonstrate the near-linear scalability of the integrated solution, it was necessary to assemble the software in a way that accurately reflects a typical scenario deployed by e-business customers. Figure 1 provides an

overview of how the software components interacted to form the benchmark environment. Quest's Benchmark Factory was used to drive the workload. For each WebSphere Application Server, Benchmark Factory simulated 1600 users making requests with a think time of 7 seconds per request. IBM's WebSphere Edge Server routed the requests to the appropriate WebSphere Application Server. WebSphere serviced the requests and interacted with the DB2 database that contained data for one million users and five thousand distinct stocks.

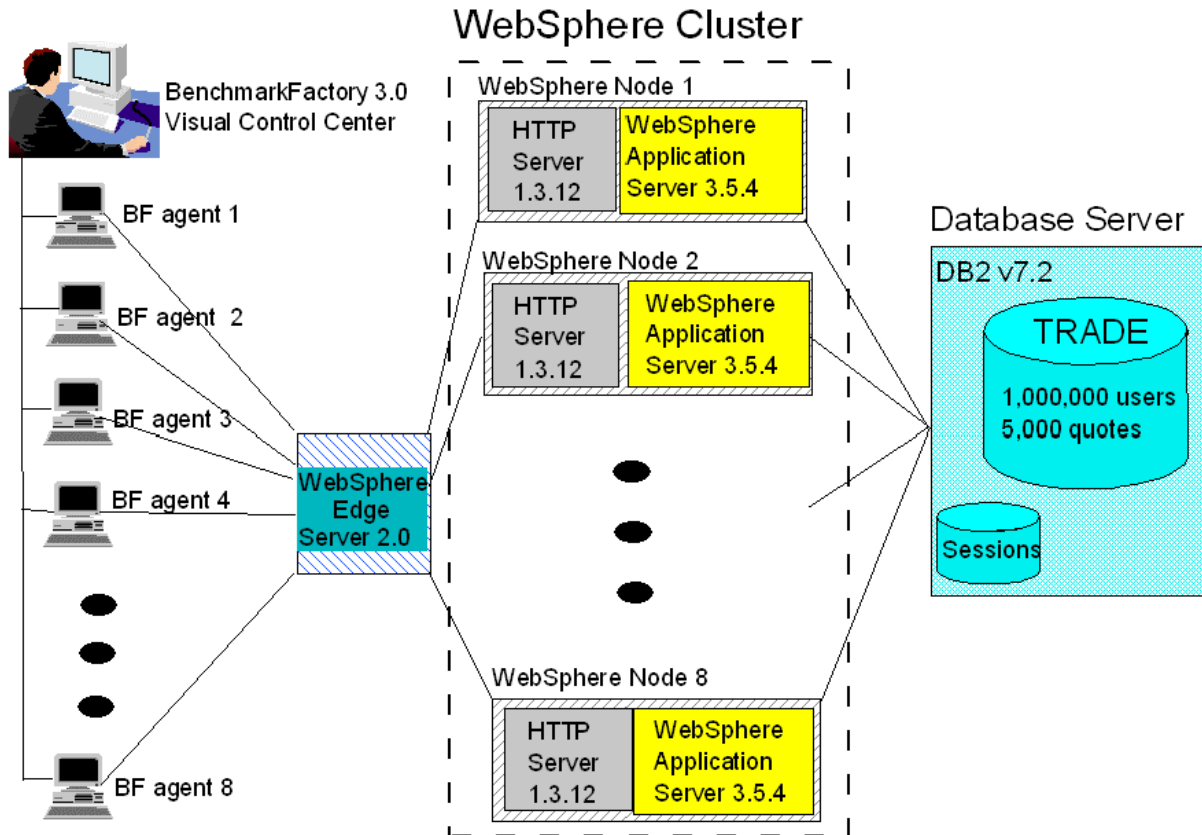


Figure 1: Software Configuration

The IBM xSeries Teraplex Integration Center provided the hardware facilities necessary to simulate the complex environment present in e-business environments. Figure 2 depicts the hardware used and how the software was deployed on those systems. The tests utilized 17 of the xSeries Model x370 servers in the Teraplex. To maintain the greatest testing flexibility, the basic configuration was not modified for most of these servers. With the exception of the DB2 server, each machine had 4 Intel Pentium III Xeon 700 MHz CPUs and 4 gigabytes of RAM. The DB2 server was upgraded to 8 CPU's. As will be seen in the final results, most of these machines were very underutilized. By maintaining a common configuration, the team was able to easily change the function of a particular machine to explore multiple configurations. The WebSphere Edge Server workload was minimal and it was run on one of the Benchmark Factory machines. All machines ran Windows 2000 Advanced Server.

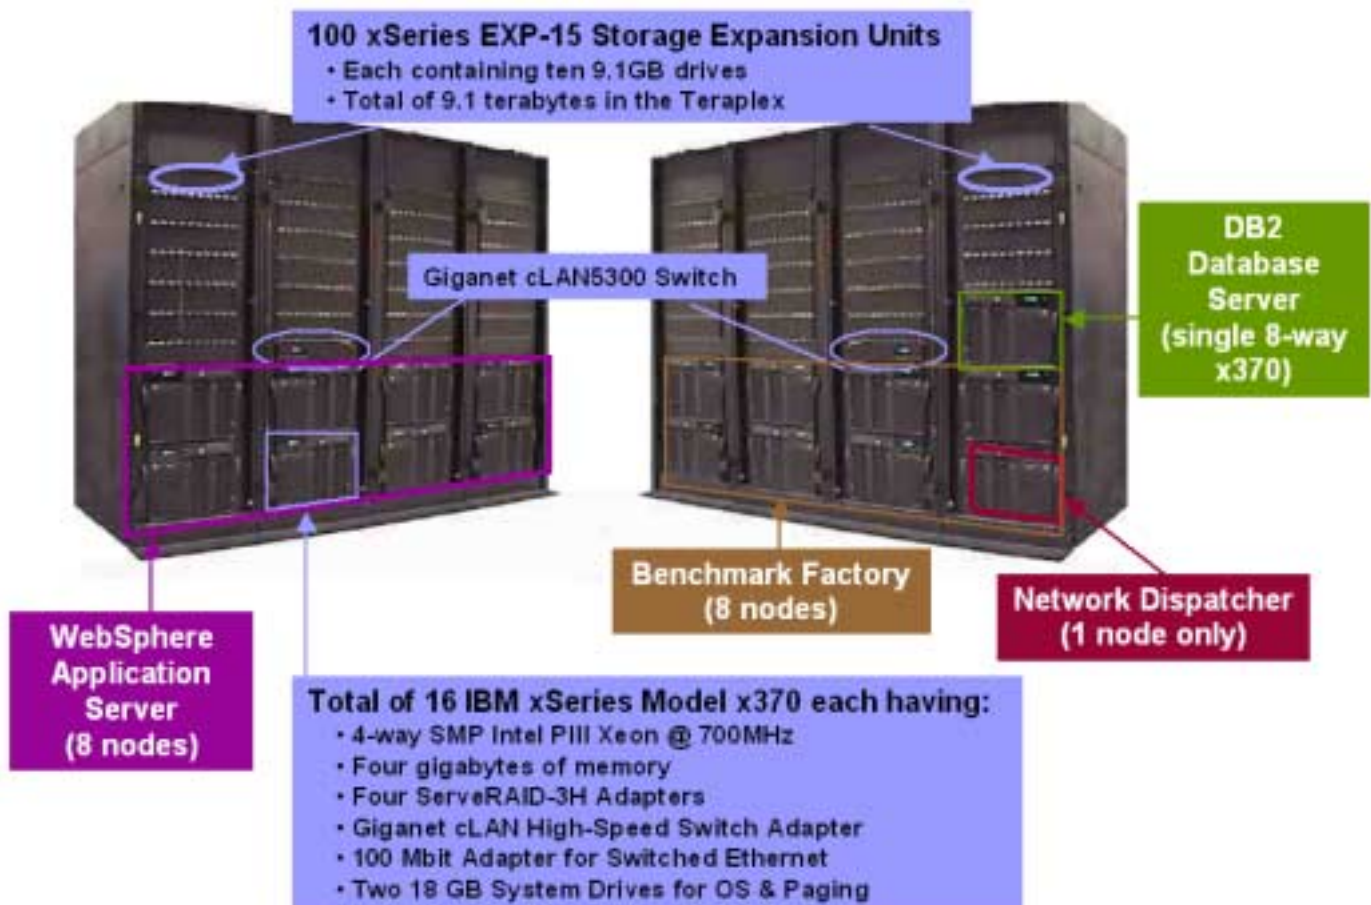


Figure 2: Hardware Configuration

Testing and Measurement Methodologies

To demonstrate the predictable scalability that is required to support the dynamic conditions prevalent in e-business, the team chose to implement the Trade2 application from the WebSphere Performance Benchmark Sample. By simulating an online brokerage firm, Trade2 provided an ideal scenario for testing the scalability of both the WebSphere Application Server and the DB2 backend. Quest also provides a module for Benchmark Factory that facilitates the automation of the Trade2 clients.

To accurately measure the scalability, the team first established a baseline using Benchmark Factory to determine the maximum number of clients and transactions per second could be achieved using a single WebSphere Application Server. Based on previous benchmarks and feedback from customer test cases, the team chose to use 7 seconds as the average think time per transaction. Given this think time, the single node tests determined that the best throughput was achieved when simulating 1600 users.

The baseline measurements were compared against results at 2, 4, 6, and 8 WebSphere nodes. As additional WebSphere nodes were added, the 1600 users per node and 7 second think time were kept constant. This allowed accurate measurement of the transactions per second as an indicator of overall system throughput. Along with

measuring the transactions per second, CPU utilization on the WebSphere nodes and DB2 server was recorded during the tests.

Along with testing overall scalability of the configuration, specific key components of the configuration were isolated for additional testing. The first such tests involve placing the HTTP session database on a separate server from the trade database to determine if any contention was introduced by hosting them on the same server. The second set of tests involved isolating which WebSphere Edge Server load balancing option would provide the best throughput as the environment scaled.

Benchmark Results

This benchmarking effort provides compelling evidence that IBM's WebSphere and DB2 provide the scalability and integration required to deliver predictable scalability for demanding e-business workloads. Figure 3 highlights the near-linear scalability obtained as the environment scaled up. Running against a single WebSphere Application Server the configuration supported 1600 users driving 216 transactions per second. As additional WebSphere nodes were added, the number of users per WebSphere Application Server were scaled linearly to enable accurate comparisons of overall transactions per second. The graph clearly shows that as the number of users and WebSphere nodes increased, the overall transactions per second also increased in a near linear fashion to handle 98,040 requests per minute from 12,800 concurrent users.

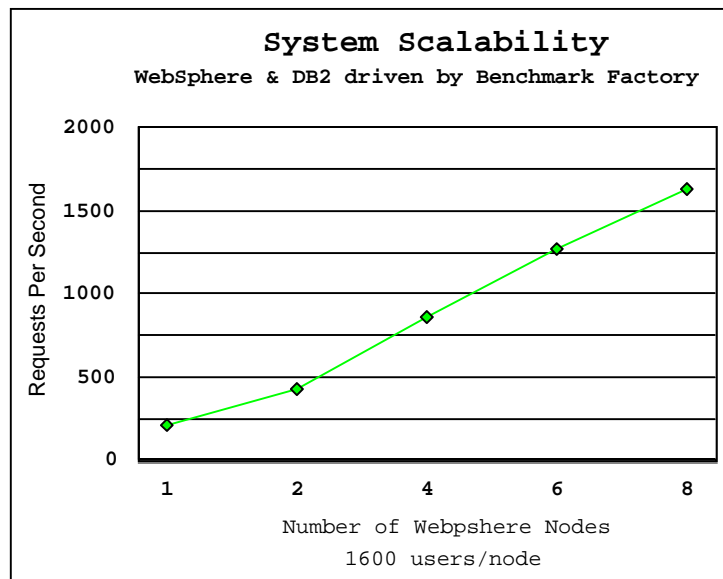


Figure 3: System Scalability

It is also important to note that even though most of the systems used in the test were configured as 4 CPU servers with 4 gigabytes of memory, this was more a matter of convenience than necessity. The machines that drove Benchmark Factory ran at less than 10% CPU utilization and the WebSphere servers ran at 85% CPU utilization.

Since the single DB2 server was a key component of the solution, the scalability of this server was also measured during the tests. Figure 4 shows how DB2 scaled to support the increasing workload in this environment by measuring CPU usage at each stage of the tests. As additional users and transactions were added, DB2 also demonstrated near-linear scalability.

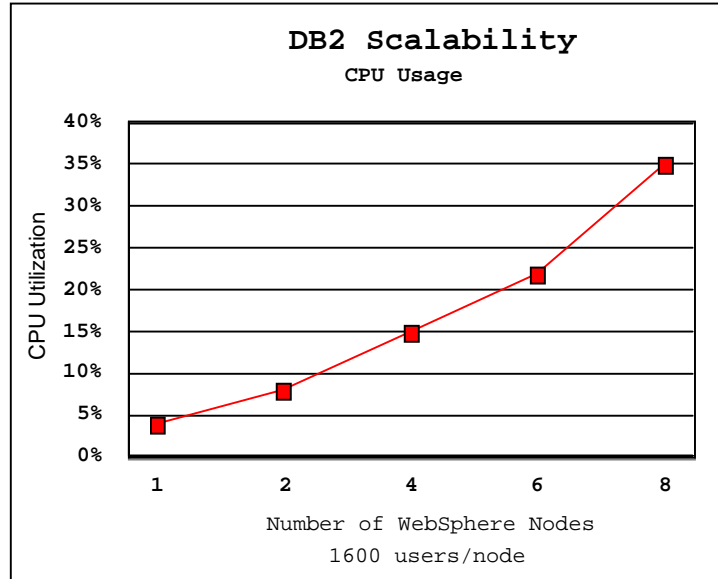


Figure 4: DB2 Scalability

At peak load, the tests showed no significant differences in overall system transactions per second when isolating the session on a separate DB2 server. This provides validation for the decision to use a single database server to support the data needs of the entire e-commerce scenario.

The WebSphere Edge Server provides a feature to load balance requests based on the IP address of the client. This load-balancing algorithm allows subsequent client requests to be directed to the same server. This feature yielded the best overall throughput in our tests.

Conclusion

Building an e-infrastructure doesn't have to be an exercise filled with uncertainty and doubt. Building on a hardware and software platform that can offer predictable levels of performance today and in the future will help alleviate many of those doubts. The results of this scalability study provides compelling evidence that WebSphere and DB2 on IBM's Intel-based servers delivers a highly scalable solution for building an infrastructure that will support demanding e-business workloads, with reliable predictability. Benchmark Factory delivers the ability to test a deployment under multiple stress factors to determine exactly what the scalability behavior will be in an e-business application.

Solution Details

Building a dynamic e-business solution like the one represented in these tests requires a partnership of technologies and skills. This section provides more details about the technologies and resources demonstrated by each of the partners in this effort.

IBM

IBM DB2 Universal Database

IBM DB2 Universal Database provides the capabilities to support e-business as well as the data integration capabilities that enable e-businesses to thrive. Robust, scalable, and reliable, DB2 delivers the foundation for building a successful e-business.

Innovation to enable a new world of possibilities

IBM invented the relational database and has been the leading force for innovation in the field ever since. That tradition continues in DB2 Universal Database Version 7.2. DB2 was the first database to deliver web services based on the XML, SOAP, and UDDI protocols enabling a whole new breed of integrated business application. Combined with additional technologies like the in-memory search capabilities provided by the DB2 Net Search Extender enable new levels of e-business performance.

More than just data management

Today's business information is no longer restricted to structured records-based documents, but encompasses text files, photographs, audio, video, spreadsheets and spatial data. Many of today's e-businesses are sitting on a gold mine of information tucked away in older legacy applications and other platforms. Ideally, this data would be seamlessly integrated with newer e-business solutions and data in a scalable, reliable and available environment. DB2 is uniquely capable of leveraging these diverse types and sources of data to generate business value. Whether it is managing non-traditional data via DB2 Extenders, integrating data from other RDBMS with Relational Connect, or integrating with other applications via XML based web services or MQ Series integration DB2 has the solution you need.

Scalability for all workloads

DB2 provides the technology required to meet the transaction and data volume requirements of the largest corporations while remaining flexible enough to serve the varying needs of small and medium-sized business. From laptops to servers to massively parallel clusters, DB2 has a proven track record of meeting the demands of all types of workloads at any scale.

Available when you need it

Today's e-businesses are always open and require databases that are always available. DB2 not only delivers the high availability features needed to keep your data always available. Beyond that, DB2 has capabilities to ensure that routine maintenance tasks don't interfere with data availability.

IBM Teraplex Integration Centers

The IBM Teraplex Integration Centers are facilities that have been put in place since 1996 to prove very large data base implementations for business intelligence solutions. IBM has invested in excess of \$76 million US dollars to build Teraplex Centers for iSeries-AS/400, xSeries-Netfinity and NUMA-Q, pSeries-RS/6000, and zSeries-S/390 server platforms. Each center is equipped with the most powerful systems of their type

With years of success testing integrated solutions integrating IBM products with non-IBM components, the Teraplex Centers also have the capabilities to address large data base e-business solutions beyond the data warehouse. Customized proofs of concept are executed using real customer data on very large-scale configurations. In addition to IBM customers, many premier software developers have also used the Teraplex Centers to help ensure their solutions function when pushed to a very large scale.

Since inception, IBM's Teraplex Centers have been bringing value to IBM customers and Business Partners with their proofs of concept to

- Minimize risk
- Prove the scalability of IBM systems
- Share skills to improve performance
- Stress test IBM Business Partner software

Find out more about the IBM Teraplex Integration Centers at
<http://www.ibm.com/software/data/bi/teraplex>

Find out more about DB2 at www.ibm.com/software/db2

IBM WebSphere Application Server

IBM WebSphere Application Server, Advanced Edition, provides a strong, cross-platform, Java technology-based Web application platform that supports deployment of e-business applications. Incorporating a high-performance Enterprise JavaBeans(EJB) server, WebSphere Application Server Advanced Edition offers extensive application management capabilities, and it supports medium- to high-level transactional environments used in conjunction with dynamic web page serving.

WebSphere Application Server Advanced Edition enables powerful interactions with and flexible connectivity to relational databases, transaction processing systems, and other applications. Also, usability enhancements across installation, administration and product documentation allow WebSphere Application Server, Advanced Edition, to accelerate the development of dynamic Web sites capable of handling advanced e-business transactions.

WebSphere Application Server V4.0¹ is leading the way in support for industry open-standards. WebSphere Application Server V4.0 provides full Java 2 Platform, Enterprise Edition (J2EE) compliance with a rich set of enterprise Java open standards implementations. It also provides built in support for the key Web services open standards, making it production-ready for the deployment of enterprise Web services solutions.

The IBM WebSphere software platform for e⁻¹business is a comprehensive set of integrated e-business solutions. It is based on industry standards that make it flexible and pluggable, which can enable you to adapt as markets shift and business goals change. Building on this robust platform, you can integrate diverse IT environments to maximize existing investments. You can deliver core business applications to the Web, and scale these applications to meet changing needs and increasing demand.

WebSphere Application Server Advanced Edition, Version 4.0 includes:

- J2EE compliance
- WebSphere editions
- Java 2 Connectors (JCA)
- Web services
- Web server plugin
- Embedded HTTP server
- Performance enhancements
- Administration tools
- Other Tools
- Expanded platform support
- Expanded database support
- Migration
- Enterprise Extensions

For more information on WebSphere Application Server, visit this IBM Web site
www.ibm.com/software/webservers/appserv

IBM xSeries Servers

The IBM family has been designed to provide the most robust and scalable foundation for today's evolving e-business demands. More than a family of servers, IBM provides an integrated approach to building a flexible e-business infrastructure using hardware, software and services.

- Scalable to accommodate your business as your capacity needs increase
- 24x7 availability to accommodate global Internet business hours

¹ WebSphere 4.0 was not generally available until after these tests were performed. WebSphere Application Server 3.5 was used for these tests.

- Designed and optimized for specific e-business workloads
- Pre-tested, best of breed application configurations
- Assurance of quality backed with support whenever you need it

IBM X-architecture™ technologies give xSeries servers their name and mainframe-like stability. In combination with the line's extraordinary management tools, these technologies set xSeries servers apart from other Intel-based systems. That competitive edge helps reduce costs across your IT infrastructure. Known as TCO (total cost of ownership), these expenses include the direct costs of technology maintenance as well as the indirect costs of lost productivity, should technical issues subject your business to unplanned outages.

Unique xSeries technologies and tools include:

- Sophisticated hot-swap components
- Proactive diagnostics and alerts
- Enhanced error-correcting memory
- Intuitive capacity management
- High-performance, data-protective storage

xSeries servers incorporate mainframe-inspired technologies and intelligent management tools for real hands-off reliability. Part of the IBM hardware portfolio for e-business, xSeries systems are a high-performance line of Intel-based servers spanning the solution spectrum. xSeries servers have three product categories:

- Point Solutions offer out-of-the-box deployment
- Universal, the most versatile of the categories, offers the ability to quickly expand or redeploy as business needs change
- Rack-optimized servers solve the data center space dilemma by delivering outstanding performance and manageability in the slimmest chassis possible. IBM engineers have perfected the sleek designs that keep these servers small, while sacrificing nothing in computing power.

Quest Software Benchmark Factory load testing software.

Benchmark Factory was selected as the workload simulator primarily for its scalability and ease of use, but also as a result of its support of the IBM Trade2 Benchmark 'out of the box'. Benchmark Factory is load testing and capacity planning software that scales throughput testing to virtually unlimited users. Benchmark Factory prevents system downtime and poor performance by determining the critical limits of any application environment before the environment is live.

Benchmark Factory exposes application reliability problems by simulating user transactions that stress all components of the environment. By simulating hundreds and thousands of complex transactions, administrators can predict system capacity limits, bottlenecks and other stress related issues within all layers of the environment. This includes determining the number of transactions the entire system can handle, as well as individual components such as database server and application server. These tests are repeatable, and can be adjusted on the fly to provide realistic scenarios that match a particular environment. Additionally, Benchmark Factory monitors a wide variety of system statistics within the tool. Benchmark Factory can also establish baselines and build detailed reports for instant analysis of results.

Benchmark Factory basics:

- **Statistical forecasting** is collecting e-commerce performance data, such as CPU, memory, disk and network use and workload. Using this historical data, future trends may be predicted. This method is similar to weather prediction –the base predictions have a reasonable level of accuracy, but unexpected variations are not uncommon. In our opinion, statistical forecasting is a good performance reality check, but should not be used as the only method of forecasting performance.
- **User simulation** is constructing e-commerce system models and applying various user loads against the system models to determine how a system responds. Constructing accurate system models is often a complex and time-

consuming process. An accurate model gives a reliable estimate of performance at various user loads, but failing to create a useful model provides little useful information.

- **Profile** is using standard benchmark routines developed by several companies and universities to test e-commerce system components. Benchmark Factory offers detailed routines for the following: AS3AP, Database Spec 94, Scalable Hardware, Set Query, TPC-B, TPC-C, TPC-D and Wisconsin. Results from these benchmarks, yield performance expectations of e-commerce hardware, but do not provide testing of the actual e-commerce application.
- **Synthetic Workload** is a hybrid of User Simulation and Profile, consisting of simulated user scenarios exercising the e-commerce application at a variety of user loads. Synthetic workload tests the actual e-commerce application and hardware and closely simulates actual user activities.

For more information on Benchmark Factory: www.quest.com/benchmark_factory/

Intel Corporation

As the world's leading supplier of processors designed for servers, Intel has developed a powerful suite of server building blocks and tools that enable original equipment manufacturers (OEM), local system integrators and resellers to build world-class servers. From server chipsets to motherboards, system chassis to RAID controllers, Intel provides the crucial building blocks for the most widely used server platforms for the Internet economy. Intel architecture based server provides both scale out and scale up capability, the scale out strategy is highly cost-effective to deploy volume servers instead of proprietary systems. Computing resources can be expanded in a modular fashion to more precisely match e-Business growth and to meet changing demands caused by the ongoing Internet and business requirements. New applications and services can also be deployed more quickly, and workloads are easily balanced across multiple systems to improve performance and availability as tested in this benchmark. This incremental approach also helps businesses take more rapid advantage of the technological innovations that are continuously happening in the horizontal server market.

- With Intel based servers you can quickly, reliably and cost-effectively scale out to add additional capacity and high availability to your e-Business data center.
- Intel based servers easily scale up with additional processors and memory for the faster response time and growing transaction volume demands of the Internet-enabled enterprise.

By scaling out and scaling up on an Intel-based server, businesses can benefit from the economies of volume, while quickly deploying the most cost-effective and efficient solutions available.

Find out more about Intel at www.intel.com/

The WebSphere Performance Benchmark Sample, a.k.a. Trade2

The WebSphere Performance Benchmark Sample, a.k.a. Trade2, was originally built by the WebSphere Performance team as a tool for testing the performance of the WebSphere Application Server. Derived from experiences with many customer environments, Trade2 is a collection of Java classes, Java Servlets, Java Server Pages and Enterprise Java Beans integrated into a single application, designed to emulate an online brokerage firm. Figure 1 shows the system topology in which the Trade2 application runs. Trade2 was developed in conformance with the “WebSphere Application Development Best Practices for Performance and Scalability” using the VisualAge for Java and WebSphere Studio tools. Each of the components is written to open web and Java Enterprise APIs, making the Trade2 application portable across J2EE compliant application servers.

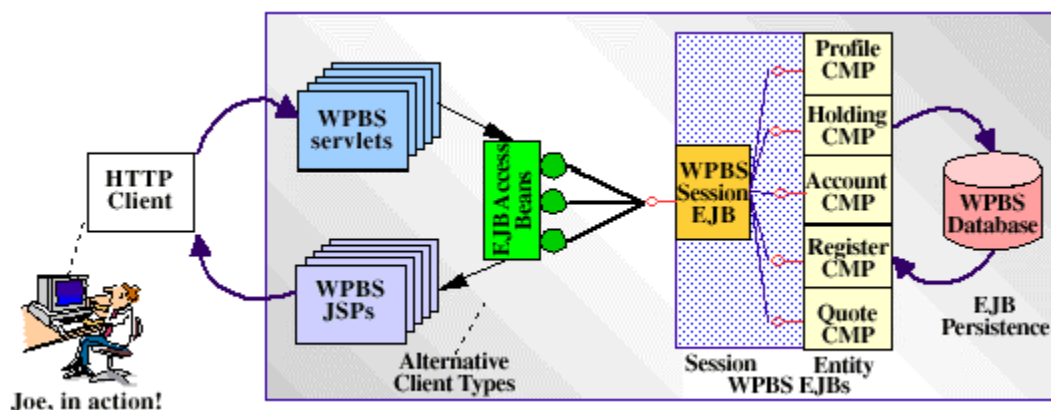


Figure 2: Trade 2 Topology

The Trade2 application allows a user, typically using a web browser, to perform the following **actions**:

- Register to create a user profile, user ID/password and initial account balance
- Login to validate an already registered user
- Browse current stock price for a ticker symbol
- Purchase shares
- Sell shares from holdings
- Browse portfolio
- Logout to terminate the users active interval

Each **action** is comprised of many primitive operations running within the context of a single HTTP request/response. For any given action there is exactly one transaction comprised of 2-5 remote method calls. A **Sell** action for example, would involve the following primitive operations:

- Browser issues an HTTP GET command on the TradeAppServlet
- TradeServlet accesses the cookie-based HTTP Session for that user
- HTML form data input is accessed to select the stock to sell
- The stock is sold by invoking the **sell()** method on the **Trade** bean, a stateless **Session EJB**. To achieve the sell, a transaction is opened and the Trade bean then calls methods on Quote, Account and Holdings **Entity EJBs** to execute the sell as a single transaction.

The results of the transaction, including the new current balance, total sell price and other data, are formatted as HTML output using a Java Server Page, portfolio.jsp.

Initially developed to assist IBM’s product development and performance test teams, Trade2 can now be downloaded from the Internet, at http://www.ibm.com/software/webservers/appserv/wpbs_download.html, by IBM customers and business partners.



© International Business Machines Corporation 2001

IBM Corporation
Marketing Communications
Software Group
Route 100
Somers, NY 10589

Printed in the United States of America.
09-01
All Rights Reserved

This publication was produced in the United States of America.
IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

IBM hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

® DB2, DB2 Universal Database, WebSphere, WebSphere Application Server, IBM, the IBM logo, and Teraplex are trademarks of International Business Machines Corporation in the United States, other countries or both.

© Intel Corporation 2001. All Rights Reserved.
*Third-party brands and names are the property of their respective owners

Intel and Xeon are trademarks or registered trademarks of Intel Corporation in the United States, other countries or both.

© Quest Software, Inc. 2001. All rights reserved.

Quest and Quest Software are trademarks, and Benchmark Factory is a registered trademark, of Quest Software, Inc. Other company, product, and service names may be trademarks or service marks of others.