IBM eServer Solutions
June 2003

Scaling Large Commercial Web-based Applications using IBM eServer Capacity on Demand

IBM eServer iSeries scales to 14,148 concurrent J.D. Edwards OneWorld Xe HTML users, processing 17 different commercial transactions using WebSphere Application Servers on IBM POWER4™ processor-based hardware

Neil Willis, Advisory Software Engineer
IBM eServer iSeries
neilwill@us.ibm.com

30 June, 2003

IBM | e business solutions | JDEDWARDS

Executive Summary

This report describes various techniques used to scale a large commercial web-based application such as J.D. Edwards OneWorld® Xe using the J.D. Edwards interactive Web/HTML benchmark running on IBM eServer iSeries servers featuring POWER4 processors.

This test successfully executed 14,148 concurrent user sessions in a multi-tiered environment. Response times were sub-second, easily meeting the J.D. Edwards certification criterion of 2 seconds.

Two different configurations were tested, to validate different customer requirements.

- A Virtual Three Tier (V3T) environment was initially scaled vertically to 11,065 users at 0.34 seconds response time
- A Physical Three Tier (P3T) environment was used to then scale horizontally to 14,148 users at 0.55 seconds.

The V3T environment used two IBM eServer iSeries model 870s, each with 16 processors and 64 GB memory connected to an IBM eServer iSeries model 890 with 32 processors with 256 GB of memory. The P3T environment used five IBM eServer iSeries model 870s, each with 16 processors and 64 GB of memory each connected to an IBM eServer iSeries model 890 with 32 processors and 256 GB of memory.

During the course of this test, the iSeries model 870s were initially shipped as 8-way systems, with the IBM Capacity on Demand feature, which allowed the number of processors to be increased to 16 without any interruption to processing.

In the J.D. Edwards benchmark environment for Xe, actual OneWorld HTML clients were used to generate a real workload to be processed by the enterprise and web servers. By doing so, this benchmark provides a realistic measure of J.D. Edwards performance and demonstrates the suitability of the IBM eServer iSeries as a superior platform for e-Business applications.
These results reflect the speed and scalability of servers based on IBM POWER4 processing technology, the advantages of a well-balanced configuration, and the ability to harness IBM eServer iSeries Capacity on Demand.

Background Information

The primary objective of the benchmark was to demonstrate the outstanding scalability of the iSeries POWER4-based server platform running OneWorld Xe SP20 B1 in a WebSphere HTML environment.

A secondary objective was to determine and document the scaling techniques used, so that customers can use similar methods to design and implement large commercial systems to cater for upwards of 10,000 web-based users.

Highlights

The Virtual 3 Tier Benchmark and the Physical 3 Tier Benchmark have eclipsed all previous J.D. Edwards benchmarks in scalability and response time. Below is a table which compares recent benchmarks, as certified by J.D. Edwards.

	Sun P3T (Feb 2003) SP17	IBM V3T (June 2001) SP14 Webp	IBM V3T (April 2003) SP20 B1	IBM V3T (June 2003) SP20 B1
Users	8,969	6,020	11,065	14,148
CPUs Technology	168 UltraSparc III 750MHz	72 ISTAR 500MHz	64 POWER4 1.3GHz	112 POWER4 1.3GHz
Users/CPU	53	84	173	126
CPU Utilization %	Web/App 67% DB 80%	Web 65% App/DB 65%	Web 62% App/DB 74%	Web/App 55% DB 51%
Resp. Time (s)	1.46	1.16	0.34	0.55
Users per WebSphere JVM	150 WebSphere 3.5	376 WebSphere 3.5	1250 WebSphere 4.0	1415 WebSphere 4.0
OS/Version DB/Version	Solaris 8 Oracle 8	OS/400 V4R5 DB2 UDB	OS/400 V5R2 DB2 UDB	OS/400 V5R2 DB2 UDB

The significant improvement in response time (from 1.16 seconds in 2001, to 0.34 seconds in 2003) can be attributed to a number of factors:

- **IBM POWER4 processor technology**: Running Java™ and C application code can benefit significantly from a faster CPU. POWER4 chips also enable the processing of multiple instructions on a single clock cycle. iSeries uses a 64-bit implementation of POWER4 technology, which allows WebSphere ap-

plications written in Java to scale better and to be more easily managed with increased I/O bandwidth. As can be seen in the chart above, the number of users per processor has increased considerably.

- **High performance disk subsystem:** The disk subsystem is based around an Ultra-wide SCSI architecture and uses PCI-X bus structure. Each subsystem uses a #2757 PCI-X Ultra RAID Disk Controller with 235MB of write cache, and fifteen 10,000 RPM 17.54 GB disk units, operating in a RAID-5 configuration. In 2001, we saw an average disk service time (controller service time plus disk service time plus queuing) of 3-5 milliseconds. This had decreased to an average 1-2 milliseconds using latest disk adapter and disk unit technology.

- **High performance and scalable HTTP server (powered by Apache)**: In this benchmark, iSeries has proven itself as a scalable and high performance web server. In the Virtual 3 Tier environment, each iSeries model 870 16-way was able to process HTML requests from over 5,500 concurrent web clients. These requests were nontrivial dynamic requests involving complex web pages and javascript. On average, the total system was handling an average of 2.9 million hits per hour, and peaking at over 3.6 million hits per hour. This workload processes 17 different transactions, so it is not simple static surfing, but more complex active interaction, similar to placing, for example, an order with 10 line items on a web site.
- **WebSphere 4.0** (with Java 1.3) has been shown to scale higher than WebSphere 3.5, conserving resources, through vertical scaling. Also the use of WebSphere cloning has improved fault tolerance. WebSphere 4.0 and Java on iSeries have been shown to perform as a very reliable platform, when stressed under heavy workloads.

- **Capacity on Demand** allows a non-disruptive way on gaining more processing power, especially in compute intensive environments such as a Java application server. This would allow customers to handle peaks in their workloads, in a cost effective way, while avoiding downtime caused by a hardware upgrade.

- **Vertical scaling using Virtual Three Tier**, gave the most efficient use of CPU power (users per CPU), while Physical Three Tier allows a module way to scale horizontally. There is an inherent overhead in using Physical Three Tier, since there are more communication links (and TCP/IP stacks) to traverse, and overall this was estimated to use an additional 15% of CPU resources. Similar overheads have been observed on other hardware platforms (however there has been no scientific evidence published to-date, other than this paper, to support this hypothesis).

J.D. Edwards n-Tier Architecture

At a high level, the architecture of the benchmark environment consisted of various tiers:

- The Database Server, accessed either through a high-level SQL interface (XDA) or through Java Data-Base Connection (JDBC).
- The OneWorld Xe Application Server which contains the J. D. Edwards business logic.
- The WebServer handles client requests from a web browser interface, such as Microsoft's Internet Explorer, using hypertext Markup Language (HTML). The WebServer uses IBM WebSphere middleware to forward HTTP requests to Java servlets which interface to the OneWorld Xe Application Server and generate HTML responses to the web client.

The concept of n-tier architecture gives considerable flexibility in the placement of each tier. In this series of benchmarks, two different configurations were used.

- Virtual Three Tier, where the OneWorld Xe Application Server is on the same system as the Database Server, and
- Physical Three Tier where the OneWorld Application Server is located on a different system to the Database Server.

Virtual Three Tier Configuration using IBM eServer iSeries

Systems

Three IBM eServer iSeries systems (two model 870 16-way systems and one model 890 32-way) were used, connected by 1 Gbps high speed fibre optics. Each CPU operated at 1.3Ghz and shared 2 MB of L2 cache. They were built using IBM copper interconnect technology and silicon-on-insulator for higher speed and lower heat dissipation.

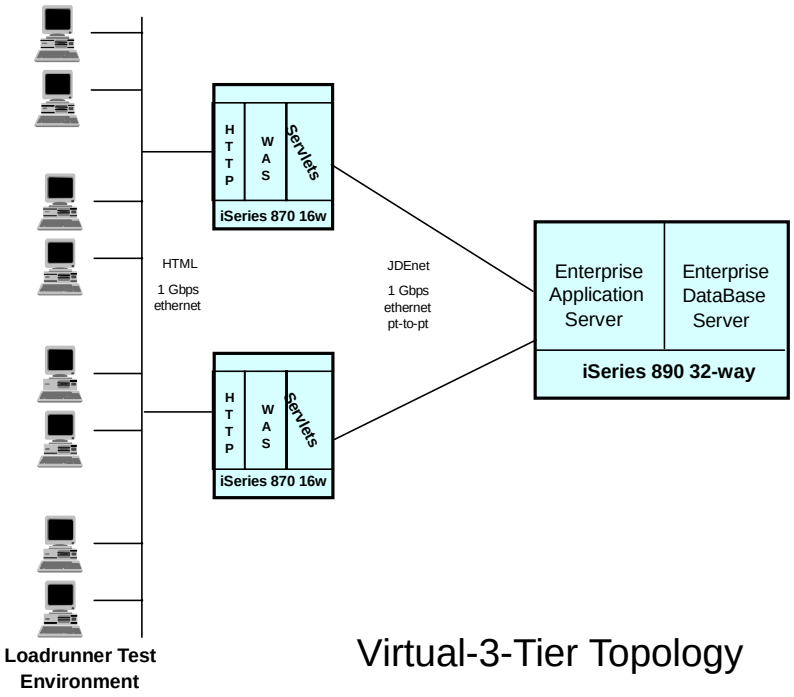

Virtual-3-Tier Topology

Database Server and Application Server: The IBM eServer iSeries model 890 with 32 processors, was configured with 256 GB of main memory and 225 disk arms in Raid-5 sets giving a useful capacity of 3.42 TB.

Webservers: Two IBM eServer iSeries model 870 were configured each with 64 GB of main memory and 90 disk arms in Raid–5 protected sets giving a useful capacity of 1.37 TB.

Software:

- OneWorld Xe SP20 B1
- WebSphere 4.0.5
- OS/400® - V5R2 Cum 4
- Mercury Loadrunner 7.51

Network: Each Web server was connected to the Enterprise server by 1 Gbps Ethernet adapters. The 1 Gbps Ethernet connections are point-to-point, full duplex over optical fiber cables. These lines had relatively light utilization (<8%), maximizing communication throughput and minimizing delay.

Loadrunner Agent PCs: (simulating web users) were connected to a Cisco router in a full duplex 1 Gbps Ethernet network.

Database
All IBM eServer iSeries models are shipped with an integrated DB2 database which was used to store all test data. The test database simulates a generic mid-to-large scale customer, matching J.D. Edwards' target market of customers with $200 million to $2 billion in annual revenues. It does not reflect any particular customer or industry, but has the flexibility to model a mixture of Distribution, Manufacturing and Financial users. The table sizes required to model a large customer were determined based on the experience of J.D. Edwards field and corporate personnel.

Physical Three Tier Configuration using IBM eServer iSeries Systems

The Physical Three-tier Configuration was the same as the Virtual Three-tier except as noted. Six IBM eServer iSeries systems (five model 870 16-way systems and one model 890 32-way) were used , connected by 1 Gbps high speed fibre optics.

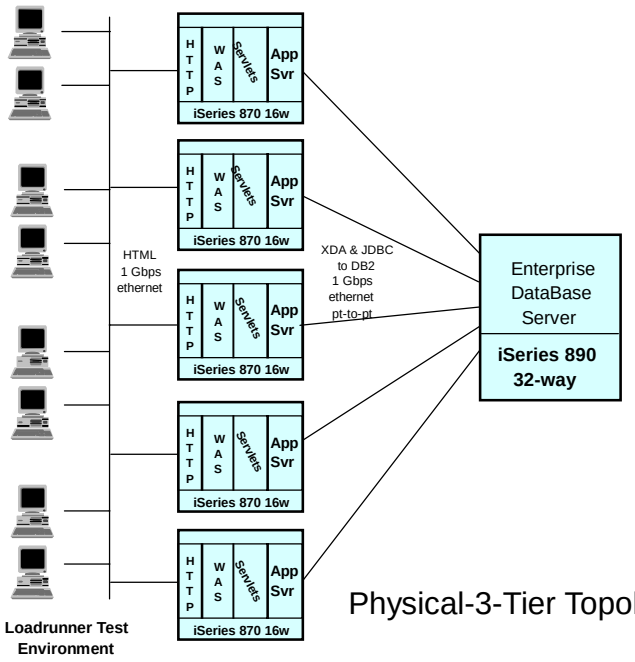

Physical-3-Tier Topology

Database Server: The IBM eServer iSeries model 890 with 32 processors, was configured with 256 GB of main memory and 225 disk arms in Raid-5 sets giving a useful capacity of 3.42 TB.

Web and Application Servers: Five (5)IBM eServer iSeries model 870s were configured each with 64 GB of main memory and 90 disk arms in Raid–5 protected sets giving a useful capacity of 1.37 TB.

Use of IBM eServer iSeries Capacity on Demand

In our initial scaling tests, the iSeries model 870s where supplied as 8-way systems with the 8/16 Capacity on Demand feature. As we began

to ramp-up, we started to over utilize the CPUs. Instead of adding more 8-way 870's, we decided to activate the Capacity on Demand features, to increase the number of available CPUs to 16.

The activation process was straightforward and took only minutes once the upgrade key was supplied by IBM. Here are the steps required:

1. Verify the target system has the Capacity on Demand feature
2. Collect Capacity on Demand information using System Service Tool (SST) utility
3. Send Capacity on Demand data to IBM
4. Receive upgrade key from IBM (assume payment terms have been met)
5. Copy key into iSeries Capacity on Demand SST screen
6. Use Partition Management (another SST screen) to add new processors

An IPL was not required, resulting in zero downtime. Capacity on Demand is available for purchase two ways – as a permanent in-crease (one time charge) or an on/off increase (paid for on a daily basis).

Vertical Scaling versus Horizontal Scaling

Many J.D. Edwards customers choose to run OneWorld Xe in a virtual three-tier configuration, with the OneWorld Xe application running on the same system as the Database server. This has certain advantages compared to running a physical three-tier configuration.

First, systems management becomes simpler, since there are less systems to manage. There is less hardware and software. Second, reliability is increased, as there are less components to fail (iSeries incorporates many features to increase reliability, such as redundant power supplies and the use of Raid-5 disk subsystems, giving it one of the highest levels of availability in the industry). Third, efficiency is greater, when running in a Virtual three-tier configuration.

In the iSeries implementation of OneWorld Xe, when running a virtual three-tier configuration, requests for data from the OneWorld Xe application server are made through a call level interface to DB2. This has the effect of using less CPU and providing a faster response time. In our tests, we used about 15% less CPU and our average response time was approximately 0.2 seconds faster.

Why then deploy physical three-tier environments? For enterprise customers who need to run large numbers of concurrent users, the Enterprise Database Server in a virtual three-tier configuration may begin to reach its limits of scalability.

To give the Enterprise Database Server more capacity, using J.D. Edwards n-tier technology, we can offload the OneWorld Xe application to another node. In the testing outlined in this paper, we actually offloaded a copy of the OneWorld Xe code onto each of the five Web server systems. By doing this we have created a convenient way of horizontal scaling. Capacity can be increased by adding more Web/Application server building blocks. The theoretical limit to scaling now becomes the Database Server and its ability handle SQL re-quests.

In summary, our recommendation is to scale vertically initially for ease of management, reliability and efficiency, but consider horizontal scaling for outright capacity.

WebSphere Cloning - Performance and Fault Tolerance

During this test, various WebSphere configurations were trialled. We wanted to not only maximize WebSphere performance but also provide a high degree a fault tolerance.

WebSphere Advanced Server Version 4 (5722-WS4) on iSeries uses a 64-bit Java Virtual Machine (JVM). This has considerable advantages over a 32-bit JVM implementation, since more objects can be addressed, and greater vertical scaling can be obtained. We found when using a 12GB initial Heap (which expands up to 20GB during our tests) that it was possible to scale a single instance of the WebSphere application up to 3,000 concurrent users.

However, to improve fault tolerance we decided to use WebSphere feature for cloning of instances. To do this, the current instance is promoted to a group, then the group can be cloned as many times as required, with each clone inheriting properties from the group. When this is done, the Apache plug-in, that directs HTML traffic to WebSphere, is able to redirect traffic to another clone, if for any reason any clones fail. We standardized on two clones, each handling 1,500 users, with initial heap sizes of 6GB. We also found that by building in this extra parallelism, we obtained slightly better response time. So on each of five iSeries 870 systems, we had two clones, giving us a theoretical maximum of 15,000 users. This could easily be increased by re-tuning our clones to a higher user count.

Be aware that creating extra instances impacts memory and CPU, since extra copies of many objects are pinned in memory. Using two clones (and VMS) per system provided us with a good balance between performance, efficiency and fault tolerance. As a footnote, it would have been possible to use multiple copies of Apache as well to enhance fault tolerance, but we did not test this environment.

Design of a High Bandwidth Network

In both virtual three-tier and physical three-tier environments, large amounts of data are transported between tiers. When designing the network, it is important to engineer sufficient bandwidth to minimize transit time and lower response time. Often for a small additional cost, routes can be replaced with point-to-point communications, which can increases data throughput, while reducing latency and eliminating some risk due to failure of communications subsystems.

Early in the project, we made a decision to use 1 gigabit per second (Gbps) ethernet for all communications. It would have been convenient to arrange for all equipment to be connected on a single router. However, based on data collected from our pilot test system, we calculated that this system would become overloaded at high user levels.

Our first decision was to introduce multiple networks, where independent data was moved to separate networks. For the HTML traffic from the Loadrunner agents (simulating thousands of browser clients), it was imperative to use a high speed router, because each Loadrunner agent PC must be able to send data to any of the WebServers.

Our second decision was to use point-to-point connections for dedicated networks from the Web Servers to the Application/Database Server or from the Web/Application Servers to the Database Server. In the schematics above, for both environments, point-to-point fiber connections were used. In both V3T and P3T environments there was no requirement for communication between the WebServers or Web/Application Servers, so effectively they have their own private network to the Application/Database Server (V3T) or Database Server (P3T).

Extra communication adapters were required on the iSeries 890 system, but there was a savings in routers. Each link runs at 1 Gbps, full duplex, with zero latency, and only transfers data relevant to that link. Using more adapters on the iSeries 890 results in more data being simultaneously transferred, since there is two 1 Gbps dedicated networks in V3T and five 1 Gbps dedicated networks in P3T. One caveat, is that the Application/Database Server (V3T) or Database Server (P3T) must have adequate bus bandwidth to support this

number of high-speed Ethernet adapters (the iSeries model 890 can support a theoretical maximum of 128 high speed adapters).
One final point is that high speed Ethernet adapters usually support a *defacto* frame size standard called “jumbo” frames, which can be up to 8996 bytes long. In the connections between iSeries systems, we used this maximum size to reduce the number of frames, and CPU required to transmit the frames. The frames transferred between iSeries systems mostly contained database row data. The older Ethernet standard of 1492 bytes for each frame would have resulted in more frames being sent and greater CPU overhead. Smaller frames were found to be adequate for communications between the agents PCs and the Web Servers.

In summary, separate the networks to increase communication bandwidth, and where prudent, use point-to-point connections and large frame sizes to increase throughput.

WebSphere, HTTP and OneWorld Xe Scaling Techniques

There are many actions which can be taken to scale OneWorld Xe on iSeries. In fact there is an IBM Redbook, which covers many of these actions in detail (see bibliography). Here are some of the key actions to take, with relation to WebSphere Application Server 4.0, Advanced Edition for iSeries (5733-WA4), HTTP Server for iSeries (powered by Apache) and OneWorld Xe SP20 B1 for iSeries.

- Make sure you are running enough threads and jobs to support the number of users, that you anticipate, to be concurrently running on the different tiers.

- In HTTP Server (powered by Apache), set the number of threads to approximately 50% of the number of concurrent users that you expect to be running. In our case it was set to 1500 threads per HTTP server.

- In WebSphere, there are a number of settings which allow more threads to start. From the Administration console, in the Web Container Service, you can set the Minimum thread size to 50% of the users per instance, the Maximum thread size to 100% of users per instance and in the HTTP transport, set maximum keep alives to 100% of the users per instance. In the Session Manager Service, change the Maximum in-memory session count to 90% of users per instance (or clone if you are using WebSphere cloning).

- WebSphere has an option in the Administrators Console to set the initial heap size for the Java Virtual Machine (JVM). Try using a value of 4 MB per user. On iSeries do not set a maximum heap value, the iSeries JVM automatically handles this. If your heap is too small you may see excessive CPU consumed by the Java Garbage Collector. If your heap is too big it may cause excessive Input/Output (I/O) to disk through page faulting.
- OW Java Application Server code uses JDBC as one of its ways of accessing DB2. It is possible to pre-start jobs which service these DB2 JDBC requests, through an iSeries CL Command. We found that a number equal to 40% of the concurrent users was sufficient, since the iSeries implementation

uses WebSphere connection pooling. Here is an example where we change the number of jobs to 1000:

```
CHGPJE SBSD(QUSRWRK) PGM(QZDASOINIT) INLJOBS(1000)
```

- OneWorld Xe has a number of kernel server jobs which handle transaction processing. It is important to have enough kernel jobs to service your transaction workload. In our testing, we used the following guidelines. We used one call object kernel job per 7.5 users, one network kernel job per 20 call object kernel jobs, and one security server for every 100 users.

The values above are offered as starting points only, and should be tested for suitability before being placed in production. More information on configuration, tuning and performance can be found in the references in the Bibliography and from J.D. Edwards Knowledge Garden at www.jdedwards.com.

Conclusion

This paper has documented the enormous steps forward in scalability of IBM eServer iSeries when running J.D. Edwards OneWorld Xe software. Advances in technology, both hardware and software, have come together to provide a powerful platform to scale commercial web based applications, while providing outstanding response time for a web based client.

While today few customers run applications with more than 10,000 concurrent users, as more applications are ported to a browser interface, the number of users will increase, often with unpredictable peaks. By demonstrating that over 14,000 concurrent web-based users can easily be handled, with room for growth, customers can feel comfortable in knowing that IBM eServer iSeries when running J.D. Edwards OneWorld Xe can scale to meet their future workloads.

Bibliography

1. J.D. Edwards OneWorld XE Implementation on IBM eServer iSeries Servers, SG24-6529-00 Redbook, published May-14-2002, www.redbooks.ibm.com
2. J.D. Edwards OneWorld Implementation for AS/400, SG24-5195-01 Redbook, published July-13-2000, www.redbooks.ibm.com
3. WebSphere 4.0 Installation and Configuration on the IBM eServer iSeries Server, SG24-6815-00 Redbook, published April-22-2002, www.redbooks.ibm.com
4. Using AS/400 Database Monitor and Visual Explain To Identify and Tune SQL Queries, REDP0502
   Redpaper, published January-9-2001, www.redbooks.ibm.com
5. iSeries Performance Capabilities Reference Guide, V5R2, SC41-0607-07 published May-2003,
   http://publib.boulder.ibm.com/iseries/v5r2/ic2924/books/sm24/as4ppcp7.pdf
6. Web Client Tuning Tips for iSeries, J.D. Edwards 5 ERP 8.0 / OneWorld Xe, SP20 & SP21 published Jan-2003, www-1.ibm.com/partnerworld/pwhome.nsf/weblook/tech_support_techdoc.html
7. IBM and J.D. Edwards Push the e-Business Envelope with OneWorld and WebSphere, REDP0431 Redpaper, published April-24-2002

About the Author

Neil Willis is an Advisory Software Engineer in ERP Development at IBM Corporation, Rochester, Minnesota. He is the author of many books and papers and has presented and taught classes worldwide. Neil has worked at IBM Rochester for 7 years in the areas of Application Development, Performance Design and ERP Development. Prior to this, he was an Advisory Systems Engineer with IBM Australia for 10 years. His current projects include high level scaling of web-based commercial software applications and unicode performance. He holds a Bachelors Degree in Computer Science, a Bachelors Degree in Electrical Engineering with Honors and a Masters Degree in Information Systems.

For more information

For additional information on integrated, collaborative enterprise solutions from IBM and J.D. Edwards, call 1 888 426-5505 or visit www.ibm-jdedwards.com/

IBM Corporation
Route 100 Somers, NY 10589 U.S.A.
ibm.com

Performance data in this document was obtained in a controlled environment with specific performance benchmarks and tools. This information is presented along with general recommendations to assist the reader to have a better understanding of IBM products. Results obtained in other environments may vary significantly and does not predict a specific customer's environment.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

The information contained in this document has not been submitted to any formal IBM test and is AS IS without warranty of any kind, express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

J.D. Edwards World Solutions Company
One Technology Way
Denver, CO 80237 U.S.A.
jdedwards.com

Printed in the United States of America

07/03

GXXX-XXXX-XX