

IBM/J.D. EDWARDS WORLD AND ONEWORLD SIZING AND PLANNING QUESTIONNAIRE

Version 23
June 2002

OneWorld and World sizing estimates are available to IBM employees, J.D.Edwards employees, authorized IBM business partners with the J.D.Edwards value add, and authorized J.D.Edwards Systems Integrators only. All others should contact their IBM representative.

Assistance with this questionnaire and OneWorld sizing or assistance with other pre-sales OneWorld configuration questions can be obtained by contacting

Eddie Quallen

E-mail: equallen@us.ibm.com

Phone: 1-415-545-3611

or

Karen Weatherby

E-mail: kmweath@us.ibm.com

Phone: 1-303-773-5091

Completed sizing questionnaires for the **Americas** should be e-mailed or faxed to:

America's IBM Techline Sizing Center

West Chester, Pennsylvania

E-mail: eSizings@us.ibm.com

Phone: **1-800-426-0222** or (770-835-6690)

Fax: **770-659 5245**

Completed sizing questionnaires for **Europe**, the **Middle East**, and **Africa**, should be submitted as follows:

IBM Personnel use web based JD Edwards sizing request generator on:

<http://w3-5.ibm.com/support/emea/techline>

IBM Business Partners use web based JD Edwards sizing request generator on:

<http://partners.boulder.ibm.com> (userid and password required for access)

All other requestors please e-mail completed questionnaire to:

JDEdwards_EMEA_Presales@uk.ibm.com or **IBMJDE@uk.ibm.com**

Completed sizing questionnaires for **Asia Pacific** regions should be submitted as follows:

ASEAN (Malaysia, India, Indonesia, Philippines, Singapore, Thailand, Vietnam, Sri Lanka)	All platforms	techadm@my.ibm.com
China	pSeries	wusj@cn.ibm.com
	iSeries, xSeries	wangsh@cn.ibm.com
Hong Kong	pSeries	wongsw@hk1.ibm.com
	iSeries	ngtck@hk1.ibm.com
Taiwan	xSeries	APtech@jp.ibm.com CC: tsuchino@jp.ibm.com
	pSeries	andrewh@tw.ibm.com
South Korea	iSeries, xSeries	erichsu@tw.ibm.com
	pSeries	hongsk@kr.ibm.com
Japan	iSeries	aysun@kr.ibm.com
	xSeries	APtech@jp.ibm.com CC: tsuchino@jp.ibm.com
	All platforms	ERPCC2@jp.ibm.com

What is a sizing estimate?

A sizing estimate is an **approximation** of the hardware resources required to support a specific software implementation. It is a pre-sales effort based on information available at a point in time, providing an **entry** into understanding the customer's hardware requirements. Customers' actual experiences will vary from the sizing estimate for many reasons, including batch and reporting workloads, and custom code. The degree of variability can range from small to very significant. IBM assumes no liability for actual results that differ from the sizing estimate.

Contents

I. Introduction	Page 3
II. Contact Information	Page 4
III. Architecture	Page 5
IV. User and OneWorld World Production Environment	Page 6
V. General system and sizing questions	Page 8
VI. Appendix – OneWorld architecture diagrams	Page 10

I. Introduction

Thank you for allowing IBM the opportunity to provide you a sizing estimate for OneWorld or World software from J.D. Edwards running on an IBM server platform. We are requesting information in order to arrive at a system recommendation and architecture based on your requirements and business environment. This information will allow us to provide you with an estimation of the processor class, memory, and disk space requirements for your database/enterprise server, and if applicable, Business Function servers, Batch servers, Windows Terminal Servers and HTML servers.

The sizing process initiated by the completion of this Sizing and Planning Questionnaire is called a pre-sales sizing estimate. A pre-sales sizing is based on minimal input and is intended to be a very rough estimate of server resource required to support a given OneWorld or World workload. The accuracy of the sizing estimate is dependent upon many variables such as the accuracy of the information provided by you, the customer, (we recognize that it may be difficult to provide exact numbers on your environment) and the actual usage of the system in a real-time environment. Sizing is an iterative process and the sizing issue should be revisited a number of times during the pre-implementation phase. Performance data gathered during the pilot phase should be used to generate a more accurate estimate of server resources required for a given OneWorld or World workload. Once you are in production, future hardware requirements should be determined through the Capacity Planning process.

In this pre-sales sizing estimate, we will provide to you an estimate of disk space required to support the OneWorld or World database based on the number of users defined in your workload estimate. This will be a very rough estimate, as the disk space required does not always correlate with the number of active users.

Currently we are able to provide pre-sales sizing estimates for a standard World implementation along with two tier implementations and three tier implementations of OneWorld. For a description of two tier and three tier implementation options please see the appendix in section 6 at the end of this document.

For World only implementations, only sections 2, 4B and 5 need to be filled in.

II. Contact Information

Company Name:	
Street Address	
City, State Zip	
Country	
Contact Person	
Phone Number	
FAX Number	
E-mail address	
Company description (distributor, mfg, etc.)	
IBM Client Rep:	
Phone Number	
E-mail Address	
IBM Product Specialist:	
Phone Number	
E-mail Address	
Business Partner Company Name:	
Contact Name	
Phone Number	
Fax Number	
E-mail address	
J. D. Edwards rep:	
Phone Number	
E-mail Address	
OMSYS Number:	

III. Architecture

A. Select database server platform (mark with an “X”)

Option	Select One
iSeries with UDB/400 database	
xSeries with DB2 UDB database	
xSeries with SQL Server database	
xSeries with Oracle database	
pSeries with DB2 UDB database	
pSeries with Oracle database	

* Sizing requests that involve more than one platform and more than one scenario may take more time to complete.

B. If sizings for other platforms are being submitted on separate documents, please indicate which additional platforms have been submitted

iSeries	
pSeries	
xSeries	

C. Select location of Business Function logic

The default IBM sizing algorithm assumes that the business application logic will run on the database server in a virtual 3-tier configuration, (unless there are “fat” clients, which run the business application logic locally). If you require a separate application server for the business logic, please specify the required platform (iSeries, xSeries or pSeries):

Platform for separate application server: _____

D. Select location of Batch processing

The default IBM sizing algorithm assumes that the batch workload will run on the database server. If you require a separate server for batch processing, please specify the required platform (iSeries, xSeries or pSeries):

Platform for separate Batch server: _____

* Typical OneWorld implementations run most batch processing on the database server.

E. Select location of Java Application Server, (WebSphere server) for Web clients (HTML or Java) (mark only one with an “X”)

Select one of the following options for HTML users	Select One
Java Application Server running on the database server, combined server	
Java Application Server on a separate, dedicated server	

The default IBM sizing algorithm assumes that the Java Application server workload will run on a separate server. Please specify the required platform for web-serving (iSeries, xSeries or pSeries):

Platform for separate Java Application Server: _____

IV. User Environment

User Counts

Please complete this information as accurately as possible. Be sure not to double count users. Classify users in the application where they spend a majority of their time each day. We are trying to size the database/enterprise server for a peak hour of a day when the most users are online and actively working.

Active users: users are judged to be active if they are clicking the “OK” button or using Fn keys 2 to 3 times a minute. Users that are signed on, but are not actually using the system should not be counted. For this sizing, the number of users reported should be the total number of active users during a peak period.

“Fat” clients: Fat clients are those that have the business application logic running locally on the client.

“Mapped” clients: Mapped clients are those that have some or all of the business application logic mapped to the enterprise server or some machine other than locally.

WTS clients: Those clients which are connected via Windows Terminal Server or Citrix.

HTML/Java clients: Those clients which use a browser interface to connect to OneWorld.

Note: Fat clients are not recommended over a Wide Area Network due to bandwidth requirements.

World/WorldVision Active Users: indicates the maximum number of active users of World or WorldVision software. Use the same definition for active user as above.

A. Define ALL OneWorld active users during a peak processing period:

Application	Active OneWorld Users			
	“Fat”	“Mapped”	WTS	HTML
Financials				
Distribution*				
Advanced Pricing				
Manufacturing				
Total				

*Divide the Distribution users between those that will use Advanced Pricing (if any) and those that will not.

B. If this will be a sizing for a coexistence OneWorld implementation (valid only for iSeries), please define the World or WorldVision user counts

Application	Active World/WorldVision Users
Financials	
Distribution	
Manufacturing	
Total	

C. Transaction Information

The standard scripts used to perform the sizing tests and benchmarks process transactions at a high rate. A typical user may not process transactions at the same rate. To assist in accurately accounting for users, please provide transaction rates below. The information required is transaction lines per hour during a peak period. If your peak period is not known then we will account for the peak by taking twice the average.

For example: If a company processes 100,000 sales orders per month, with an average of 5 lines per order. Given a 20 workday month and a 10 hour work day. We would use 100,000 orders x 5 lines per order / 20 days / 10 hours per day = 2500 lines per hour average. To size for a peak period, we would assume twice the average or 5000 lines per hour.

Peak Sales Orders per hour: _____
Average Lines per sales Order: _____

***Only transactions entered by users should be included. Transactions entered via EDI will be accounted for below.**

D. EDI:

Will your environment be receiving transactions through EDI?	Yes/No
Approximate peak lines per hour to be processed in each time frame	Day _____ Night _____

E. Batch production environment

A standard IBM sizing accounts for a significant batch workload, and fits most environments. In certain cases, there are mission critical batch processes that need to be processed at the highest priority. The default assumption if no input is provided to this question is that some capacity for typical batch jobs to execute is reserved, but all critical batch processes will be run overnight or during non-peak interactive users periods. **Care should be taken when answering this question since each specified performance critical batch process during the same time period will require an additional processor for the server running the batch processes.**

How many performance critical concurrent batch jobs (UBEs) will you need to run while your interactive OneWorld users are active?	
--	--

V. General system and sizing questions

A. Additional information for production environment:

Are you currently running OneWorld or World in production?	OneWorld __ (y/n) World ____ (y/n)
Is coexistence with World software a requirement?	Yes/No
What hardware do you have now? Specify model number if known.	
Date for start of production	
Growth projected one year after start of production (%)?	
Are Multiple Instances of OneWorld required? If yes, how many?	Yes/No # of instances ____
Will IBM ESS (Enterprise Storage Server) or FAStT (Fibre Array Storage Technology), external disk be used for database storage requirements?	Yes/No
Will this be Direct Attached or in a SAN (Storage Area Network)?	Yes/No
High Availability configuration required such as MIMIX, OMS, Datamirror, HACMP, or MSCS?	MIMIX OMS Datamirror HACMP MSCS

Explanations of some of the above questions:

EDI

Electronic Data Interchange is a method of importing data such as sales orders through a batch process. Documents are transferred from vendors or other sources and converted to OneWorld transactions via the import process. Transaction volumes for sizing purposes are converted to lines per hour, such as sales order lines per hour, which need to be processed during a peak period.

Batch Requirements

UBEs (Universal Batch Engines) are OneWorld batch and reporting processes. Most OneWorld batch processes are CPU intensive. Because of this, running multiple concurrent UBEs on a single processor will cause some degradation in performance of the UBE. Since users generally process at a higher priority, their performance should not be impacted as much. In most environments, this UBE degradation will be acceptable and running more than 1 batch process per processor will provide adequate performance. In some environments there are critical batch processes that must complete in the shortest time possible. In these cases, the system is sized to allow for an extra processor for each concurrent performance critical UBE. Care should be taken in answering the number of critical batch processes. Multiple concurrent critical UBEs will require multiple processors and can quickly drive the cost of the system higher. A standard IBM sizing accounts for a significant batch workload and should suffice in most environments.

Multiple Instances of OneWorld

Some circumstances require multiple instances of the OneWorld application running on the same installation. Implementing multiple OneWorld instances requires duplicates of the all OneWorld objects, including database files and application code. Some examples of when multiple OneWorld instances may be required include support for multiple double byte languages or multiple code pages. When multiple instances are required careful customized planning for handling consolidated reporting, merged or separate file schemes, languages and time zones is needed. Additional disk space must be added to allow

for the duplication of files for each instance. Additional servers, such as TSE servers may also be required. To determine if multiple instances are required and the number of instances necessary, contact your J.D.Edwards representative.

FAStT

FAStT500 Storage Server is the storage of choice for those wanting the ultimate in midrange storage consolidation and data sharing on multiple or heterogeneous server platforms. The FAStT500 supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for business-critical applications.

The people of IBM Storage Systems can provide flexible business solutions that maximize price-to-performance ratios. IBM TotalStorage™ FAStT500 Storage Server is a full-fibre, storage area network (SAN) solution that gives you the performance you need at the right price.

ESS

The Enterprise Storage Server (ESS) is IBM's next generation leadership disk subsystem. ESS represents the integration of multiple advanced technologies for safely storing, quickly accessing, and easily managing data. The ESS can be described as an all-in-one disk system for the entire enterprise that stands apart from the competition in many ways.

Delivering innovations in design, performance, and function, the ESS meets and exceeds the stringent requirements that demanding customers place on large scale disk subsystems. Satisfying needs ranging from easier management of storage resources, to 24x365 access, to higher levels of performance and lower cost of ownership. The ESS provides demonstrable business value. With concurrent support for a broad range of server families such as Windows NT/2000 servers, UNIX based servers, and iSeries (AS/400) servers (as well as zSeries) - the ESS is truly a storage solution for the entire enterprise.

High Availability

A highly available system can be implemented using technologies including IBM's HACMP product on the UNIX platform, MIMIX (from Lakeview Corp.), OMS (from Vision Solutions), or Datamirror (from Datamirror Corp.) on the iSeries, or MSCS on the xSeries. In all cases, a second system is required for implementation.

VI. Appendix – sizing scenarios explanation

Currently we are able to provide pre-sales sizing estimates for two tier implementations and three tier implementations of OneWorld. In a two tier implementation, the proposed server functions as the central database server supporting client PCs running OneWorld business logic and the OneWorld user interface. A two-tier implementation of OneWorld is shown below:

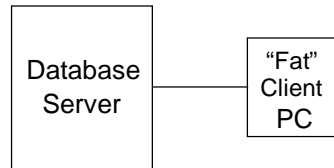


Diagram 1. 2-tier

In a three-tier implementation, the business logic (called a Business Function) can be moved from the client PC to the database server, this is known as a virtual three-tier implementation. Alternatively, the Business Function logic can be mapped to an intermediate layer on another system, this is known as a physical three-tier implementation. These two options are shown here:

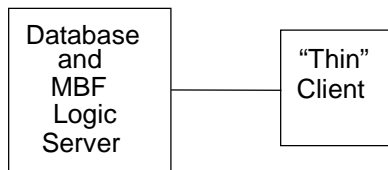


Diagram 2. Virtual 3-tier

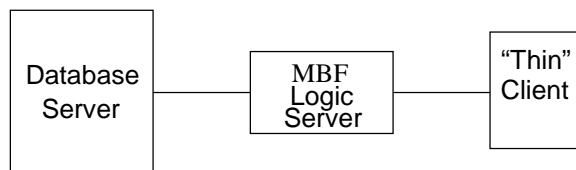


Diagram 3. Physical 3-tier

Support of WAN attached clients with OneWorld is best achieved by using Windows Terminal Server technology or by using the OneWorld HTML / Java client support. These options are shown below:

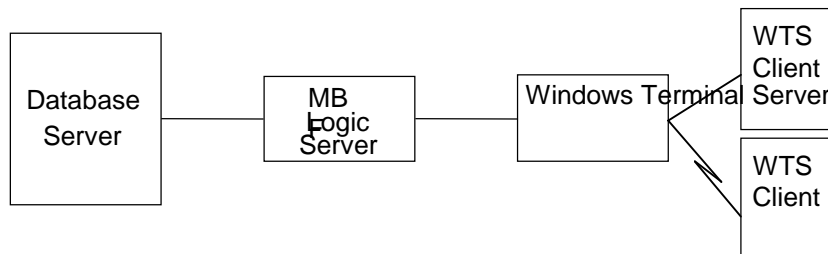


Diagram 4. Physical 3-tier with Windows Terminal Server

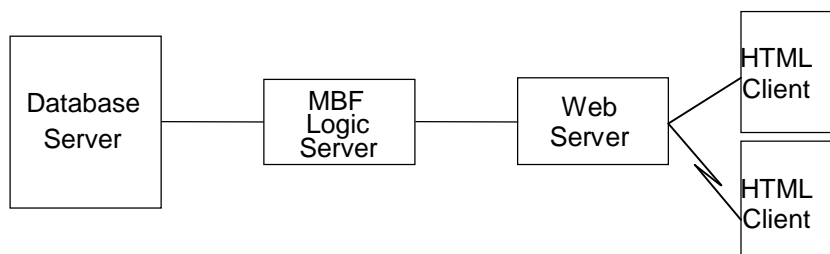


Diagram 5. Physical 3-tier with HTML / Java Clients