AS/400 Performance Capabilities Reference Version 4, Release 4

August 1999



This document is intended for use by qualified performance related programmers or analysts from IBM, IBM Business Partners and IBM customers using AS/400e series. Information in this document may be readily shared with IBM AS/400 customers to understand the performance and tuning factors in OS/400 Version 4 Release 4. Requests for use of performance information by the technical trade press or consultants should be directed to Systems Performance Department V3T, IBM Rochester Lab, in Rochester, MN. USA.

Note!

Before using this information, be sure to read the general information under "Special Notices."

Thirteen Edition (Aug 1999)

This edition applies to Version 4, Release 4 of the AS/400 Operating System.

You can request a copy of this document by download from AS/400 On Line Library via the AS/400 Internet site at: http://www.as400.ibm.com. The Version 3 Release 3 Performance Capabilities Guide is also available on the IBM AS/400 Internet site in the "On Line Library", at: http://www.as400.ibm.com. Both documents are viewable/downloadable in Adobe Acrobat (.pdf) format. Approximately 1.5MB download. Adobe Acrobate reader plug-in is available at: http://www.adobe.com.

To request the CISC version (V3R2 and earlier), enter the following command on VM:

REQUEST V3R2 FROM FIELDSIT AT RCHVMW2 (your name

To request the IBM AS/400 Advanced 36 version, enter the following command on VM:

TOOLCAT MKTTOOLS GET AS4ADV36 PACKAGE

8 Copyright International Business Machines Corporation 1999. All rights reserved. Note to U.S. Government Users -- Documentation related to restricted rights -- Use, duplication, or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

_

Special Notices	. 7
Purpose of this Document	. 9
Related Publications / Documents	. 9
Chapter 1. Introduction	10
Chapter 2. AS/400 RISC Server Model Performance Behavior	11
2.1 Overview	11
2.2 Server Model Behavior	12
2.3 Server Model Differences	14
2.4 Performance Highlights of New Model 7xx Servers	16
2.5 Performance Highlights of Current Model 170 Servers	17
2.6 Performance Highlights of Custom Server Models	18
2.7 Additional Server Considerations	18
2.8 Interactive Utilization	19
2.9 Server Dynamic Tuning (SDT)	20
2.10 Managing Interactive Capacity	23
2.11 Migration from Traditional Models	26
2.12 Migration from Server Models	28
2.13 AS/400e Dedicated Server for Domino Performance Behavior	29
Chapter 3. Batch Performance	37
3.1 Effect of CPU speed on Batch	37
3.2 Effect of DASD type on Batch	37
3.3 Tuning Parameters for Batch	38
3.4 V4R4 comments	39
Chapter 4. DB2 UDB for AS/400 Performance	40
4.1 V4R4 Enhancements for DB2 UDB for AS/400	40
4.2 Previous Version 4 Enhancements for DB2 for AS/400	43
4.3 Version 3 DB2 for AS/400 Performance Information	49
Chapter 5. Communications Performance	77
5.1 TCP/IP, Sockets, SSL, VPN, and FTP Performance Information	77
5.2 APPC, ICF, CPI-C, and Anynet Performance Information	81
5.3 LAN and WAN Performance Information	83
5.4 Work Station Connectivity Performance Information	87
5.5 Opti-Connect for OS/400 Performance Information	90
5.6 NetPerf Workload Description	92
Chapter 6. Web Serving Performance	93
6.1 Web Serving with the HTTP Server	94
6.2 Net.Commerce Performance	100
6.3 Firewall Performance	101
Chapter 7. Java Performance	102
7.1 Improved Computational Performance	102
7.2 Reduced Main Storage Footprint	104
7.3 JDBC Improvements and Commercial Java Applications	105
7.4 Unmatched Scalability	106
7.5 Comparison to Existing Languages	106
7.6 Java Performance Tips and Techniques	108
/./ Capacity Planning and Model Recommendations	115

The March Of Time	117
Better Machines for Java	117
Machines to Be Used With Caution with Java	117
Machines to Be Used With Extreme Caution for Java	118
Chapter 8. IBM Network Station Performance	120
8.1 IBM Network Station Network Data	120
8.2 IBM Network Station Initialization	121
8.3 AS/400 5250 Applications	129
8.4 Browser	129
8.5 Java Virtual Machine Applets/Applications	129
8.6 The AS/400 as a Router	129
8.7 Conclusions	130
Chapter 9. AS/400 File Serving Performance	133
9.1 AS/400 File Serving Performance	133
9.2 AS/400 NetServer File Serving Performance	133
Chapter 10. DB2/400 Client/Server and Remote Access Performance	135
10.1 Client Performance Comparisons	135
10.2 AS/400 Toolbox for Java	136
10.3 Client Access/400	138
10.4 Tips for Improving C/S Performance	147
Chapter 11. Domino for the AS/400	156
11.1 Workload Description:	156
11.2 Domino for AS/400, Release 5.0.	156
11.3 V4R4 changes that affected Domino Mail Serving Performance:	158
11.4 AS/400e Dedicated Server for Domino:	159
11.5 Mail Serving Performance Conclusions/Recommendations:	160
11.6 Domino Performance Tips/Techniques:	160
11.7 Domino Subsystem Tuning	163
11.8 Mail Serving Capacity Planning	164
11.9 Mail Serving Performance Measurements:	164
Chapter 12. OS/400 Integration of Lotus Notes Performance	167
12.1 Number of Notes Clients Supported	168
12.2 Lotus Notes DB2 Integration Performance	178
Chapter 13. Language Performance	185
13.1 Compile Time Performance	187
13.2 Runtime Performance	190
13.3 Program Object Size Comparisons	191
13.4 Working Memory Guidelines for Compiles	192
13.5 Application Compute Intensive Performance	193
Chapter 14. DASD Performance	196
14.1 Device Performance Characteristics	196
14.2 DASD Performance - Interactive	200
14.3 DASD Performance - Batch	212
14.4 DASD Performance - General	215
14.5 Integrated Hardware Disk Compression (IHDC)	216
14.6 DASD Subsystem Performance Improvements for V4R4	229
Chapter 15. Save/Restore Performance	243
15.1 Use Optimum Block Size (USEOPTBLK) parameter	244
15.2 Data Compression (DTACPR)	244

15.3 Data Compaction (COMPACT)	244
15.4 Work Loads	245
15.5 Comparing Performance Data	246
15.6 Lower Performing Tape Drives	247
15.7 Medium Performing Tape Drives	247
15.8 Highest Performing Tape Drives	247
15.9 Multiple Tape Drives	248
15.10 V4R3 Save and Restore Rates	251
15.11 V4R4 Rates On New Devices	251
15.12 Save/Restore Rates for Optical Device	252
15.13 Hierarchical Storage Management	252
15.14 Save/Restore Tips for Better Performance	253
15.15 New For V4R4	254
Chapter 16 IPL Performance	255
16.1 IPL Performance Considerations	255
16.2 IPL Benchmark Description	255
Large System Benchmark Information	256
Small System Benchmark Information	256
16.3 IPL Performance Measurements	257
16.4 MSD Affects on IPL Performance Measurements	258
16.5 IPL Tips	258
Chapter 17. Integrated Netfinity Server	259
17.1 Introduction	259
17.2 NT Server Benchmark: NetBench 5.01	260
17.3 Effects of NetBench on the AS/400	262
17.4 Performance Tips and Techniques	264
Optimizing NT Server	264
Key items to monitor on Integrated Netfinity Server.	266
Key items to monitor on AS/400	266
Component ReportComponent Interval ActivityNetBench run 200 MHz IPCS	267
17.5 Summary	267
17.6 NetBench Benchmark Details	268
17.7 Additional Sources of Information	270
Chapter 18. Logical Partitioning (LPAR)	271
18.1 Introduction	271
18.2 Considerations	271
18.3 Performance on a 12-way system	272
18.4 LPAR Measurements	275
18.5 Summary	276
Chapter 19. Miscellaneous Performance Information	277
19.1 Public Benchmarks (TPC-C, SAP, RPMark, NotesBench)	277
19.2 Dynamic Priority Scheduling	278
19.3 Main Storage Sizing Guidelines	281
19.4 Memory Tuning Updates	282
19.5 User Pool Faulting Guidelines	283
19.6 Cryptography Performance	285
19.7 AS/400 NetFinity Capacity Planning	288
Chapter 20. General Performance Tips and Techniques	291
20.1 Adjusting Your Performance Tuning for Threads	291

<i>History</i>
The Coming Change
<i>Problem</i>
<i>Solution</i>
Appendix A. CPW Benchmark Description
Appendix B. AS/400 Sizing
B.1 IBM Workload Estimator for AS/400
B.2 AS/400 CAPACITY PLANNER (BEST/1 for the AS/400) 300
B.3 BATCH400
Appendix C. DASD IOP Device Characteristics
Appendix D. AS/400 CPW Values
D.1 V4R4 Additions
D.2 AS/400e Model 7xx Servers
D.3 Model 170 Servers
D.4 AS/400e Servers
D.5 AS/400e Custom Servers
D.6 AS/400 Advanced Servers
D.7 AS/400e Custom Application Server Model SB1
D.8 Previous AS/400 RISC System Capacities
D.9 AS/400 CISC Model Capacities

Special Notices

DISCLAIMER NOTICE

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 distributed AS IS. 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.

The following terms, which may or may not be denoted by an asterisk (*) in this publication, are trademarks of the IBM Corporation.

- · · · · · ·		
AS/400	System/370	Operating System/400
C/400	AS/400e	AS/400e series
OS/400	COBOL/400	Application System/400
PS/2	RPG/400	OfficeVision
OS/2	CallPath	Facsimile Support/400
DB2	DRDA	Distributed Relational Database Architecture
AFP	SQL/400	Advanced Function Printing
IBM	ImagePlus	Operational Assistant
SQL/DS	VTAM	Client Series
400	APPN	Workstation Remote IPL/400
CICS	SystemView	Advanced Peer-to-Peer Networking
S/370	ValuePoint	OfficeVision/400
RPG IV	DB2/400	AS/400 Advanced Application Architecture
AIX	ADSM/400	ADSTAR Distributed Storage Manager/400
IPDS	AnyNet/400	IBM Network Station

The following terms, which may or may not be denoted by a double asterisk (**) in this publication, are trademarks or registered trademarks of other companies as follows:

TPC Benchmark	Transaction Processing Performance Council
TPC-A, TPC-B	Transaction Processing Performance Council
TPC-C, TPC-D	Transaction Processing Performance Council
Lotus Notes, Lotus, Word Pro	Lotus Development Corporation
Notes, 123, CC Mail, Freelance	Lotus Development Corporation
Microsoft, Windows 95	Microsoft Corporation
Windows 95, Windows 95 Explorer	Microsoft Corporation
Microsoft Word, PowerPoint, Excel	Microsoft Corporation
ODBC, Windows NT Server, Access	Microsoft Corporation
Visual Basic, Visual C++	Microsoft Corporation
Adobe PageMaker	Adobe Systems Incorporated
Borland Paradox	Borland International Incorporated
CorelDRAW!	Corel Corporation
dBASEIII Plus	Borland International
Paradox	Borland International
WordPerfect	Satelite Software International
BEST/1	BGS Systems, Inc.
NetWare	Novell
Compaq	Compaq Computer Corporation
Proliant	Compaq Computer Corporation
BAPCo	Business Application Performance Corporation
Harvard	Gaphics Software Publishing Corporation
HP-UX	Hewlett Packard Corporation
HP 9000	Hewlett Packard Corporation
INTERSOLV	Intersolve, Inc.
Q+E	Intersolve, Inc.
Netware	Novell, Inc.
Pentium	Intel Corporation
SPEC	Syems Performance Evaluation Cooperative
UNIX	UNIX Systems Laboratories
WordPerfect	WordPerfect Corporation
Powerbuilder	Powersoft Corporation
SQLWindows	Gupta Corporation
NetBench	Ziff-Davis Publishing Company
DEC Alpha	Digital Equipment Corporation
Java	Sun Microsystems, Inc.

Other terms that are used in this document may be trademarks of other companies.

V4R4 Performance Capabilities Reference 8 Copyright IBM Corp. 1999

Purpose of this Document

The intent of this document is to provide guidance in terms of AS/400 performance, capacity planning information, and tips to obtain best performance. This document is typically updated with each new release or more often if needed. This August 1999 edition of the V4R4 Performance Capabilities Reference Guide is an update to the May 1999 edition to reflect new product functions announce on August 03, 1999, with available through the remainder of 1999. The August 1999 edition superceeds the May 1999 edition.

The V4R3 version of this document, dated September 1998, includes updates since the V4R2 version, dated February 1998.

For performance information on AS/400 Upgrade to RISC, refer to the document *AS/400 Upgrade Timing Information* that is available on MKTTOOLS. The package name is AS400TIM.

The wide variety of applications available makes it extremely difficult to describe a "typical" workload. The data in this document is the result of measuring or modeling certain application programs in very specific and unique configurations, and should not be used to predict specific performance for other applications. The performance of other applications can be predicted using a system sizing tool such as BEST/1(**) for OS/400 (refer to AXC for more details on BEST/1 support).

Related Publications / Documents

The following publications/documents are considered particularly suitable for additional information on AS/400 performance topics.

- AS/400 Programming: Work Management Guide, SC41-4306
- AS/400 Programming: Performance Tools/400 Guide, SC41-8084
- AS/400 Performance Management V3R1, GG24-3723-02
- The presentation AS/400 Versus Microsoft's SNA Server Gateway

This presentation and script discusses the advantages and disadvantages of the SNA Server approach. Included are the results of an IBM benchmark that compares the SNA Server to IBM direct attachment. To receive this package type the following on the PROFS command line:

TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET AS4VSSNA PACKAGE

If you are unable to access the information indicated above, please contact your IBM technical representative.

Chapter 1. Introduction

V4R4 continues to enhance the AS/400e series value proposition - the best melding of a superior operating system with 64-bit RISC hardware. The performance of V4R4 is greatly improved by both software enhancements and new hardware. AS/400e series continues to deliver customer usable performance in the multi-user, multi-applications environment by supporting interactive, client server, batch, groupware (Domino), Java, business intelligence, and web (e-business) serving.

High-end 8 and 12-way processors were enhanced in V4R3, increasing performance of AS/400e series by 94% more than that available in V4R2. V4R3 and V4R4 supports new Model 720, 730 and 740 servers, new Model 170s and new Dedicated Servers for Domino. The servers offer up to 100% more performance than original V4R3 models. In addition, the new models offer interactive performance features for the new server/interactive processing flexibility. These new models offer better interactive/server mode performance and significant price/performance improvements. There is now over a 300-fold range of scalability in performance from the smallest e-series model to the largest 12-way.

The primary V4R4 performance items are:

- New models, both entry point, middle range and high-end growth
- Excellent Lotus Domino mail performance and new Dedicated Domino Servers for even better price/performance
- Improved V4R4 CPW values, faster processors, more memory, larger L2 cache and disk capacity
- New higher-performance and higher capacity I/O tape, disk (up to 4.2TB) and optical
- Enhanced server model performance characteristics and interactive/server algorithm
- Improved NT performance on Integrated Netfinity® Server for AS/400
- Reduced storage cost with possible performance improvements using first-to-market integrated hardware data compression (IHDC) for both disk and tape compression
- Support for 10k rpm disks, 1.6GB Read Cache and IBM Versatile Storage Server for sharing disks
- Logical Partioning (LPAR) of OS/400 for multiple simultaneous partions with separate processors, storage, clock, primary language and currency capabilities
- New Universal Database support for new image, video, audio and other larger object types in DB2/400
- Enhanced support for continous availability clustering
- New secure Enterprise-class TCP/IP support and dramatic TCP/IP performace improvements
- Dramatic server-side Java performance measured by Business Object Benchmark for Java (jBOB)
- Parallel save/restore performance improvements with hierarchical storage management support based on user-defined policies and parallel single object support for multiple tapes
- Parallel data load and index maintenance for faster databases
- Improved Query performance on n-ways with encoded vector indexes
- Web server and secure web server performance improvements

Customers who wish to remain with their existing hardware but want to move to the V4R4 operating system may find functional and performance improvements. In some cases however, N-way systems that are limited due to high CPU utilization or high contention in V4R3 may see a small increase in CPU % utilization with V4R4. This increase (less than 8%) is required to support new millenium functions like the hypervisor, teraspace, large object data types, etc.

Version 4 Release 4 OS/400 continues to protect the customer's investment while providing more function, growth, capacity, performance and better price/performance over previous versions.

Chapter 2. AS/400 RISC Server Model Performance Behavior

2.1 Overview

AS/400* Advanced Servers and AS/400e* servers are intended for use primarily in client/server or other non-interactive work environments. 5250-based interactive work can be run on these servers, but with limitations.

AS/400e Dedicated Server for Domino models will be generally available on September 24, 1999. Please refer to Section 2.13, *AS/400e Dedicated Server for Domino Performance Behavior*, for additional information.

The underlying performance structure of the AS/400e custom servers is the same as that of the AS/400 Advanced Servers and the AS/400e servers. AS/400e custom servers are designed to provide optimally balanced performance in non-interactive and interactive environments as required by specific ISV software. The primary environment addressed is one which runs significant interactive 5250 applications and is migrating toward a client/server "graphical" version of the 5250 applications as well as new client/server applications. Depending on the interactive feature installed, the new model 7xx servers are designed for non-interactive work and a limited amount of 5250-based interactive activity.

For the rest of this chapter, server will generically be used to represent AS/400e servers, AS/400e custom servers, AS/400 Advanced Servers, and 7xx model servers.

5250-based interactive work (hereafter called interactive) is defined as any job doing 5250 display device I/O. This includes:

All 5250 sessions	RUMBA/400
Any green screen interface	Screen scrapers
Telnet or 5250 DSPT workstations	Interactive subsystems
5250/HTML workstation gateway	Twinax printer jobs
PC's using 5250 emulation	BSC 3270 emulation
Interactive program debugging	5250 emulation
PC Support/400 work station function	

In general, the Type column of WRKACTJOB will tell you which jobs are considered interactive (type = INT). These are all jobs that were initiated by signing on at a 5250 display device.

Another general way to determine interactive work is to run the Performance Monitor (STRPFRMON, ENDPFRMON) during selected time periods and review the first page of the Performance Tools/400 licensed program Component Report. Additional details on monitoring interactive activity are provided later in the section, **2.10 Managing Interactive Capacity**.

Non-interactive work (hereafter called client/server) is defined to be any application that is not interactive such as batch, client/server, database, etc.

The performance information and equations in this paper represent ideal environments. This information is presented along with general recommendations to assist the reader to have a better understanding of the AS/400 server models. Actual results may vary significantly.

This chapter is organized into the following sections:

- Server Model Behavior
- Server Model Differences
- Performance Highlights of New Model 7xx Servers
- Performance Highlights of Current Model 170 Servers
- Performance Highlights of Custom Server Models
- Additional Server Considerations
- Interactive Utilization
- Server Dynamic Tuning (SDT)
- Managing Interactive Capacity
- Migration from Traditional Models
- Migration from Server Models
- AS/400e Dedicated Server for Domino Performance Behavior

2.2 Server Model Behavior

Server model behavior applies to:

- AS/400 Advanced Servers
- AS/400e servers
- AS/400e custom servers
- AS/400e model 150
- AS/400e model 170
- AS/400e model 7xx

Relative performance measurements are derived from commercial processing workload (CPW) on AS/400. CPW is representative of commercial applications, particularly those that do significant database processing in conjunction with journaling and commitment control.

Traditional (non-server) AS/400 system models had a single CPW value which represented the maximum workload that can be applied to that model. This CPW value was applicable to either an interactive workload, a client/server workload, or a combination of the two.

Now there are two CPW values. The larger value represents the maximum workload the model could support if the workload were entirely client/server (i.e. no interactive components). This CPW value is for the processor feature of the system. The smaller CPW value represents the maximum workload the model could support if the workload were entirely interactive. For 7xx models this is the CPW value for the interactive feature of the system.

The two CPW values are NOT additive - interactive processing will reduce the system's client/server processing capability. When 100% of client/server CPW is being used, there is no CPU available for interactive workloads. When 100% of interactive capacity is being used, there is no CPU available for client/server workloads.

For model 170s announced in 9/98 and all subsequent systems, the published interactive CPW represents the point (the "knee of the curve") where the interactive utilization may cause increased overhead on the system. (As will be discussed later, this threshold point (or knee) is at a different value for previously announced server models.) Up to the knee the server/batch capacity is equal to the processor capacity (CPW) minus the interactive workload. As interactive requirements grow beyond the knee, overhead grows at a rate which can eventually eliminate server/batch capacity and limit additional interactive growth. It is best for interactive workloads to execute below (less than) the knee of the curve. (However, for those models having the knee at 1/3 of the total interactive capacity, satisfactory performance can be achieved.) The following graph illustrates these points.



Figure 2.1. Server Model behavior

The figure above shows a straight line for the effective interactive utilization. Real/customer environments will produce a curved line since most environments will be dynamic, due to job initiation, interrupts, etc.

In general, a single interactive job will not cause a significant impact to client/server performance

Microcode task CFINTn, for all AS/400 models, handles interrupts, task switching, and other similar system overhead functions. For the server models, when interactive processing exceeds a threshold amount, the additional overhead required will be manifest in the CFINTn task. Note that a single interactive job will not incur this overhead.

There is one CFINTn task for each processor. For example, on a single processor system only CFINT1 will appear. On an 8-way processor, system tasks CFINT1 through CFINT8 will appear. It is possible to

see significant CFINT activity even when server/interactive overhead does not exist. For example if there are lots of synchronous or communication I/O or many jobs with many task switches.

The effective interactive utilization (EIU) for a server system can be defined as the useable interactive utilization plus the total of CFINT utilization.

2.3 Server Model Differences

Server models were designed for a client/server workload and to accommodate an interactive workload. When the interactive workload exceeds an interactive CPW threshold (the "knee of the curve") the client/server processing performance of the system becomes increasingly impacted at an accelerating rate beyond the knee as interactive workload continues to build. Once the interactive workload reaches the maximum interactive CPW value, all the CPU cycles are being used and there is no capacity available for handling client/server tasks.

Custom server models interact with batch and interactive workloads similar to the server models but the degree of interaction and priority of workloads follows a different algorithm and hence the knee of the curve for workload interaction is at a different point which offers a much higher interactive workload capability compared to the standard server models.

For the server models the knee of the curve is approximately:

- 100% of interactive CPW for:
 - AS/400e model 170s announced on or after 9/98
 - 7xx models
- 6/7 (86%) of interactive CPW for:
 - AS/400e custom servers
- 1/3 of interactive CPW for:
 - AS/400 Advanced Servers
 - AS/400e servers
 - AS/400e model 150
 - AS/400e model 170s announced in 2/98

For the 7xx models the interactive capacity is a feature that can be sized and purchased like any other feature of the system (i.e. disk, memory, communication lines, etc.).

The following charts show the CPU distribution vs. interactive utilization for Custom Server and pre-2/99 Server models.



Figure 2.2. Custom Server Model behavior



Figure 2.3. Server Model behavior

2.4 Performance Highlights of New Model 7xx Servers

7xx models were designed to accommodate a mixture of traditional "green screen" applications and more intensive "server" environments. Interactive features may be upgraded if additional interactive capacity is required. This is similar to disk, memory, or other features.

Each system is rated with a **processor CPW** which represents the relative performance (maximum capacity) of a processor feature running a commercial processing workload (CPW) in a client/server environment. **Processor CPW** is achievable when the commercial workload is not constrained by main storage or DASD.

Each system may have one of several interactive features. Each interactive feature has an **interactive CPW** associated with it. **Interactive CPW** represents the relative performance available to perform host-centric (5250) workloads. The amount of interactive capacity consumed will reduce the available processor capacity by the same amount. The following example will illustrate this performance capacity interplay:



Figure 2.4. Model 7xx Utilization Example

At 110% of percent of the published interactive CPU, or 32.1% of total CPU, CFINT will use an additional 39.8% (approximate) of the total CPU, yielding an effective interactive CPU utilization of approximately 71.9%. This leaves approximately 28.1% of the total CPU available for client/server work. Note that the CPU is completely utilized once the interactive workload reaches about 34%. (CFINT would use approximately 66% CPU). At this saturation point, there is no CPU available for client/server.

2.5 Performance Highlights of Current Model 170 Servers

AS/400e Dedicated Server for Domino models will be generally available on September 24, 1999. Please refer to Section 2.13, *AS/400e Dedicated Server for Domino Performance Behavior*, for additional information.

Model 170 servers (features 2289, 2290, 2291, 2292, 2385, 2386 and 2388) are significantly more powerful than the previous Model 170s announced in Feb. '98. They have a faster processor (262MHz vs. 125MHz) and more main memory (up to 3.5GB vs. 1.0GB). In addition, the interactive workload balancing algorithm has been improved to provide a linear relationship between the client/server (batch) and published interactive workloads as measured by CPW.

The CPW rating for the maximum client/server workload now reflects the relative processor capacity rather than the "system capacity" and therefore there is no need to state a "constrained performance" CPW. This is because some workloads will be able to run at processor capacity if they are not DASD, memory, or otherwise limited.

Just like the model 7xx, the current model 170s have a **processor capacity** (CPW) value and an interactive capacity (CPW) value. These values behave in the same manner as described in the **Performance highlights of new model 7xx servers** section.

As interactive workload is added to the current model 170 servers, the remaining available client/server (batch) capacity available is calculated as: **CPW (C/S batch) = CPW(processor) - CPW(interactive)** This is valid up to the published interactive CPW rating. As long as the interactive CPW workload does not exceed the published interactive value, then interactive performance and client/server (batch) workloads will be both be optimized for best performance. **Bottom line, customers can use the entire interactive capacity with no impacts to client/server (batch) workload response times.**

On the current model 170s, if the **published interactive capacity** is exceeded, system overhead grows very quickly, and the client/server (batch) capacity is quickly reduced and becomes zero once the interactive workload reaches 7/6 of the published interactive CPW for that model.

The absolute limit for dedicated interactive capacity on the current models can be computed by multiplying the published interactive CPW rating by a factor of 7/6. The absolute limit for dedicated client/server (batch) is the published processor capacity value. This assumes that sufficient disk and memory as well as other system resources are available to fit the needs of the customer's programs, etc. Customer workloads that would require more than 10 disk arms for optimum performance should not be expected to give optimum performance on the model 170, as 10 disk access arms is the maximum configuration.

When the model 170 servers are running less than the published interactive workload, no Server Dynamic Tuning (SDT) is necessary to achieve balanced performance between interactive and client/server (batch) workloads. However, as with previous server models, a system value (QDYNPTYADJ - Server Dynamic Tuning) is available to determine how the server will react to work requests when interactive workload exceeds the "knee". If the QDYNPTYADJ value is turned on, client/server work is favored over additional interactive work. If it is turned off, additional interactive work is allowed at the expense of low-priority client/server work. QDYNPTYADJ only affects the server when interactive requirements exceed the published interactive capacity rating. The shipped default value is for QDYNPTYADJ to be turned on.



The next chart shows the performance capacity of the current and previous Model 170 servers.

Figure 2.5. Previous vs. Current Server 170 Performance

2.6 Performance Highlights of Custom Server Models

Custom server models were available in releases V4R1 through V4R3. They interact with batch and interactive workloads similar to the server models but the degree of interaction and priority of workloads is different, and the knee of the curve for workload interaction is at a different point. When the interactive workload exceeds approximately 6/7 of the maximum interactive CPW (the knee of the curve), the client/server processing performance of the system becomes increasingly impacted. Once the interactive workload reaches the maximum interactive CPW value, all the CPU cycles are being used and there is no capacity available for handling client/server tasks.

2.7 Additional Server Considerations

It is recommended that the System Operator job run at runpty(9) or less. This is because the possibility exists that runaway interactive jobs will force server/interactive overhead to their maximum. At this point it is difficult to initiate a new job and one would need to be able to work with jobs to hold or cancel runaway jobs.

You should monitor the interactive activity closely. To do this take advantage of PM/400 or else run the Performance Monitor tool nearly continuously and query monitor data base each day for high interactive

use and higher than normal CFINT values. The goal is to avoid exceeding the threshold (knee of the curve) value of interactive capacity.

2.8 Interactive Utilization

When the interactive CPW utilization is beyond the knee of the curve, the following formulas can be used to determine the effective interactive utilization or the available/remaining client/server CPW. *These equations apply to all server models*.

```
CPWcs(maximum)= client/server CPW maximum valueCPWint(maximum)= interactive CPW maximum valueCPWint(knee)= interactive CPW at the knee of the curveCPWint= interactive CPW of the workload
```

 ${f X}$ is the ratio that says how far into the overhead zone the workload has extended:

X = (CPWint - CPWint(knee)) / (CPWint(maximum) - CPWint(knee))

EIU = Effective interactive utilization. In other words, the free running, **CPWint(knee)**, interactive plus the combination of interactive and overhead generated by **X**.

```
EIU = CPWint(knee) + (X * (CPWcs(maximum) - CPWint(knee)))
```

CPW remaining for batch = CPWcs(maximum) - EIU

Example 1:

A model 7xx server has a Processor CPW of **240** and an Interactive CPW of **70**. The interactive CPU percent at the knee equals (70 CPW / 240 CPW) or **29.2%**. The maximum interactive CPU percent (7/6 of the Interactive CPW) equals (81.7 CPW / 240 CPW) or **34%**.

Now if the interactive CPU is held to less than **29.2**% CPU (the knee), then the CPU available for the System, Batch, and Client/Server work is **100% - the Interactive CPU used**.

If the interactive CPU is allowed to grow above the knee, say for example **32.1** % (110% of the knee), then the CPU percent remaining for the Batch and System is calculated using the formulas above:

X = (32.1 - 29.2) / (34 - 29.2) = .604 EIU = 29.2 + (.604 * (100 - 29.2)) = 71.9%

CPW remaining for batch = 100 - 71.9 = 28.1%

Note that a swing of + or -1% interactive CPU yields a swing of effective interactive utilization (EIU) from 57% to 87%. Also note that on custom servers and 7xx models, environments that go beyond the interactive knee may experience erratic behavior.

Example 2:

A Server Model has a Client/Server CPW of **450** and an Interactive CPW of **50**. The maximum interactive CPU percent equals (50 CPW / 450 CPW) or **11%**. The interactive CPU percent at the knee is 1/3 the maximum interactive value. This would equal **4%**. Now if the interactive CPU is held to less than 4% CPU (the knee), then the CPU available for the System, Batch, and Client/Server work is **100% - the Interactive CPU used**.

If the interactive CPU is allowed to grow above the knee, say for example 9% (or 41 CPW), then the CPU percent remaining for the Batch and System is calculated using the formulas above: X = (9 - 4) / (11 - 4) = .71 (percent into the overhead area)

X = (9 - 4) / (11 - 4) = .71 (percent EIU = 4 + (.71 * (100 - 4)) = 72%

CPW remaining for batch = 100 - 72 = 28%

Note that a swing of + or -1% interactive CPU yields a swing of effective interactive utilization (EIU) from 58% to 86%.

On earlier server models, the dynamics of the interactive workload beyond the knee is not as abrupt, but because there is typically less relative interactive capacity the overhead can still cause inconsistency in response times.

2.9 Server Dynamic Tuning (SDT)

Logic was added in V4R1 and is still in use today so customers could better control the impact of interactive work on their client/server performance. Note that with the new Model 170 servers (features 2289, 2290, 2291, 2292, 2385, 2386 and 2388) this logic only affects the server when interactive requirements exceed the published interactive capacity rating. For further details see the section, **Performance highlights of current model 170 servers**.

Through dynamic prioritization, all interactive jobs will be put lower in the priority queue, approximately at the knee of the curve. Placing the interactive jobs at a lesser priority causes the interactive jobs to slow down, and more processing power to be allocated to the client/server processing. As the interactive jobs receive less processing time, their impact on client/server processing will be lessened. When the interactive jobs are no longer impacting client/server jobs, their priority will dynamically be raised again.

The dynamic prioritization acts as a regulator which can help reduce the impact to client/server processing when additional interactive workload is placed on the system. In most cases, this results in better overall throughput when operating in a mixed client/server and interactive environment, but it can cause a noticeable slowdown in interactive response.

To fully enable SDT, customers MUST use a non-interactive job run priority (RUNPTY parameter) value of 35 or less (which raises the priority, closer to the default priority of 20 for interactive jobs).

Changing the existing non-interactive job's run priority can be done either through the Change Job (CHGJOB) command or by changing the RUNPTY value of the Class Description object used by the non-interactive job. This includes IBM-supplied or application provided class descriptions.

Examples of IBM-supplied class descriptions with a run priority value higher than 35 include QBATCH and QSNADS and QSYSCLS50. Customers should consider changing the RUNPTY value for QBATCH and QSNADS class descriptions or changing subsystem routing entries to not use class descriptions QBATCH, QSNADS, or QSYSCLS50.

If customers modify an IBM-supplied class description, they are responsible for ensuring the priority value is 35 or less after each new release or cumulative PTF package has been installed. One way to do this is to include the Change Class (CHGCLS) command in the system Start Up program.

NOTE: Several IBM-supplied class descriptions already have RUNPTY values of 35 or less. In these cases no user action is required. One example of this is class description QPWFSERVER with RUNPTY(20). This class description is used by Client Access database server jobs QZDAINIT (APPC) and QZDASOINIT (TCP/IP).

The system deprioritizes jobs according to groups or "bands" of RUNPTY values. For example, 10-16 is band 1, 17-22 is band 2, 23-35 is band 3, and so on.

Interactive jobs with priorities 10-16 are an exception case with V4R1. Their priorities will not be adjusted by SDT. These jobs will always run at their specified 10-16 priority.

When only a single interactive job is running, it will not be dynamically reprioritized.

When the interactive workload exceeds the knee of the curve, the priority of all interactive jobs is decreased one priority band, as defined by the Dynamic Priority Scheduler, every 15 seconds. If needed, the priority will be decreased to the 52-89 band. Then, if/when the interactive CPW work load falls below the knee, each interactive job's priority will gradually be reset to its starting value when the job is dispatched.

If the priority of non-interactive jobs are not set to 35 or lower, SDT stills works, but its effectiveness is greatly reduced, resulting in server behavior more like V3R6 and V3R7. That is, once the knee is exceeded, interactive priority is automatically decreased. Assuming non-interactive is set at priority 50, interactive could eventually get decreased to the 52-89 priority band. At this point, the processor is slowed and interactive and non-interactive are running at about the same priority. (There is little priority difference between 47-51 band and the 52-89 band.) If the Dynamic Priority Scheduler is turned off, SDT is also turned off.

Note that even with SDT, the underlying server behavior is unchanged. Customers get no more CPU cycles for either interactive or non-interactive jobs. SDT simply tries to regulate interactive jobs once they exceed the knee of the curve.

Obviously systems can still easily exceed the knee and stay above it, by having a large number of interactive jobs, by setting the priority of interactive jobs in the 10-16 range, by having a small client/server workload with a modest interactive workload, etc. The entire server behavior is a partnership with customers to give non-interactive jobs the bulk of the CPU while not entirely shutting out interactive.

To enable the Server Dynamic Tuning enhancement ensure the following system values are on: (the shipped defaults are that they are set on)

- QDYNPTYSCD this improves the job scheduling based on job impact on the system.
- QDYNPTYADJ this uses the scheduling tool to shift interactive priorities after the threshold is reached.

The Server Dynamic Tuning enhancement is most effective if the batch and client/server priorities are in the range of 20 to 35.

Server Dynamic Tuning Recommendations

On the new systems and mixed mode servers have the QDYNPTYSCD and QDYNPTYADJ system value set on. This preserves non-interactive capacities and the interactive response times will be dynamic beyond the knee regardless of the setting. Also set non-interactive class run priorities to less than 35.

On earlier servers and 2/98 model 170 systems use your interactive requirements to determine the settings. For "pure interactive" environments turn the QDYNPTYADJ system value off. in mixed environments with important non-interactive work, leave the values on and change the run priority of important non-interactive work to be less than 35.



Figure 2.6.

2.10 Managing Interactive Capacity

Interactive/Server characteristics in the real world.

Graphs and formulas listed thus far work perfectly, provided the workload on the system is highly regular and steady in nature. Of course, very few systems have workloads like that. The more typical case is a dynamic combination of transaction types, user activity, and batch activity. There may very well be cases where the interactive activity exceeds the documented limits of the interactive capacity, yet decreases quickly enough so as not to seriously affect the response times for the rest of the workload. On the other hand, there may also be some intense transactions that force the interactive activity to exceed the documented limits interactive feature for a period of time even though the average CPU utilization appears to be less than these documented limits.

For 7xx systems, current 170 systems, and mixed-mode servers, a goal should be set to only rarely exceed the threshold value for interactive utilization. This will deliver the most consistent performance for both interactive and non-interactive work.

The questions that need to be answered are:

- 1. "How do I know whether my system is approaching the interactive limits or not?"
- 2. "What is viewed as 'interactive' by the system?"
- 3. "How close to the threshold can a system get without disrupting performance?"

This section attempts to answer these questions.

Observing Interactive CPU utilization

The most commonly available method for observing interactive utilization is the Performance Monitor used in conjunction with the Performance Tools program product. The monitor collects data for each job on the system, including the CPU consumed and the type of job. By examining the reports generated by the Performance Tools product, or by writing a query against the data in the QAPMJOBS file (or the QAPMJOBL file in V4R4 and beyond).

The following query will yield the information you need:

Select SUM(JBCPU)/(SUM(INTSEC) *1000) as CPUPERCENT from QAPMJOBS where JBTYPE = "I".

However, this will only show an average interactive utilization for the duration of a measurement interval (Smallest is 5 minutes, default is 15 minutes). Also, as will be described later in this section, the utilizations listed for job-type "I" are not necessarily all "interactive".

There are other means for determining interactive utilization more precisely. The easiest of these is the performance monitoring function of Management Central, which became available with V4R3. Management Central can provide:

- Graphical, real-time monitoring of interactive CPU utilization
- Creation of an alert threshold when an alert feature is turned on and the graph is highlighted
- Creation of an reverse threshold below which the highlights are turned off
- Multiple methods of handling the alert, from a console message to the execution of a command to the forwarding of the alert to another system.

By taking the ratio of the Interactive CPW rating and the Processor CPW rating for a system, one can determine at what CPU percentage the threshold is reached (This ratio works for the 7xx models and the current model 170 systems. For earlier models, refer to other sections of this document to determine what fraction of the Interactive CPW rating to use.) Depending on the workload, an alert can be set at some percentage of this level to send a warning that it may be time to redistribute the workload or to consider upgrading the interactive feature.

Another method is to combine the information provided by the WRKSYSACT command and the performance monitor. The WRKSYSACT command will give a fairly accurate accounting of work being done by each task in the system for intervals of 5 seconds or greater (A larger value is recommended for balancing the impact of the command with the work on the system, although it has the advantage over the performance monitor in that it only looks at active jobs, so it does not need to page in information about all jobs). The performance monitor database can determine which jobs are listed as interactive (In V4R3, use JBTYPE = "I" in the QAPMJOBS file. In V4R4, a more accurate determination can be made by examining JBSTSF = 1 in the QAPMJOBOS file. A join query between the file generated by WRKSYSACT (QAITMON) and the QAPMJXx file can give a fairly good picture of what the interactive utilization was when the measurement was taken.

With V4R4, the new Performance Collection functions that are available can yield similar results without having to run both the monitor functions and WRKSYSACT. The collection services functions can break the data down into very small time-slices (15 seconds), so the QAPMJOBOS file can be queried directly.

Finally, the functions of PM400 can also show the same type of data that the Performance Monitor shows, with the advantage of maintaining a historical view, and the disadvantage of being only historical. However, signing up for the PM400 service can yield a benefit in determining the trends of how interactive capacities are used on the system and whether more capacity may be needed in the future.

Is Interactive really Interactive?

Earlier in this document, the types of jobs that are classified as interactive were listed. In general, these jobs all have the characteristic that they have a 5250 workstation communications path somewhere within the job. It may be a 5250 data stream that is translated into html, or sent to a PC for graphical display, but the work on the AS/400 is fundamentally the same as if it were communicating with a real 5250-type display. However, there are cases where jobs of type "I" may be charged with a significant amount of work that is not "interactive". Some examples follow:

- Job initialization: If a substantial amount of processing is done by an interactive job's initial program, prior to actually sending and receiving a display screen as a part of the job, that processing may not be included as a part of the interactive work on the system. However, this may be somewhat rare, since most interactive jobs will not have long-running initial programs.
- More common will be parallel activities that are done on behalf of an interactive job but are not done within the job. There are two database-related activities where this may be the case.
 - 1. If the QQRYDEGREE system value is adjusted to allow for parallelism or the CHGQRYA command is used to adjust it for a single job, queries may be run in service jobs which are not interactive in nature, and which do not affect the total interactive utilization of the system. However, the work done by these service jobs is charged back to the interactive job. In this case, the performance monitor and most other mechanisms will all show a higher sum of interactive CPU

utilization than actually occurs. The exception to this is the WRKSYSACT command, which will show both the current activity for the service jobs and the activity that they have "charged back" to the requesting jobs. Thus, in this situation it is possible for WRKSYSACT to show a lower system CPU utilization that the sum of the CPU consumption for all the jobs.

2. A similar effect can be found with index builds. If parallelism is enabled, index creation (CRTLF, Create Index, Open a file with MAINT(*REBUILD), or running a query that requires an index to be build) will be sent to service jobs that operate in non-interactive mode, but charge their work back to the job that requested the service. Again, the work does not count as "interactive", but the performance data will show the resource consumption as if they were.

There are two key ideas in the statements above. First, if the workload has a significant component that is related to queries, it will be possible to show an interactive utilization in the performance tools that is significantly higher than what would be assumed from the ratings of the Interactive Feature and the Processor Feature. Second, although it may make monitoring interactive utilization slightly more difficult, in the case where the workload has a significant query component, it may be beneficial to set the QQRYDEGREE system value to allow at least 2 processes, so that index builds and many queries can be run in non-interactive mode. Of course, if the nature of the query is such that it cannot be split into multiple tasks, the whole query is run inside the interactive job, regardless of how the system value is set.

How close to the threshold can a system get without disrupting performance?

The answer depends on the dynamics of the workload, the percentage of work that is in queries, and the projected growth rate. It also may depend on the number of processors and the overall capacity of the interactive feature installed. For example, a job that absorbs a substantial amount of interactive CPU on a uniprocessor may easily exceed the threshold, even though the "normal" work on the system is well under it. On the other hand, the same job on a 12-way can use at most 1/12th of the CPU, or 8.3%. a single, intense transaction may exceed the limit for a short duration on a small system without adverse affects, but on a larger system the chances of having multiple intense transactions may be greater.

With all these possibilities, how much of the Interactive feature can be used safely? A good starting point is to keep the average utilization below about 70% of the threshold value (Use double the threshold value for the servers and earlier Model 170 systems that use the 1/3 algorithm described earlier in this document.) If the measurement mechanism averages the utilization over a 15 minute or longer period, or if the workload has a lot of peaks and valleys, it might be worthwhile to choose a target that is lower than 70%. If the measurement mechanism is closer to real-time, such as with Management Central, and if the workload is relatively constant, it may be possible to safely go above this mark. Also, with large interactive features on fairly large processors, it may be possible to safely go to a higher point, because the introduction of workload dynamics will have a smaller effect on more powerful systems.

As with any capacity-related feature, the best answer will be to regularly monitor the activity on the system and watch for trends that may require an upgrade in the future. If the workload averages 60% of the interactive feature with almost no overhead, but when observed at 65% of the feature capacity it shows some limited amount of overhead, that is a clear indication that a feature upgrade may be required. This will be confirmed as the workload grows to a higher value, but the proof point will be in having the historical data to show the trend of the workload.

2.11 Migration from Traditional Models

This section describes a suggested methodology to determine which server model is appropriate to contain the interactive workload of a traditional model when a migration of a workload is occurring. It is assumed that the server model will have both interactive and client/server workloads.

To get the same performance and response time, from a CPU perspective, the interactive CPU utilization of the current traditional model must be known. Traditional CPU utilization can be determined in a number of ways. One way is to sum up the CPU utilization for interactive jobs shown on the Work with Active Jobs (WRKACTJOB) command.

Work with Active Jobs

10/22/97

CPU%: 33.0 Elapsed time: 00:00:00 Active jobs: 152

Type options, press Enter.

2=Change 3=Hold 4=End 5=Work with 6=Release 7=Display message 8=Work with spooled files 13=Disconnect ...

Opt	Subsystem/Job	User	Туре	CPU %	6 Function	Status
	BATCH	QSYS	SBS	0		DEQW
	QCMN	QSYS	SBS	0		DEQW
	QCTL	QSYS	SBS	0		DEQW
	QSYSSCD	QPGMR	BCH	0	PGM-QEZSCNEP	EVTW
	QINTER	QSYS	SBS	0		DEQW
	DSP05	TESTER	INT	0.2	PGM-BUPMENUNE	DSPW
	QPADEV0021	TEST01	INT	0.7	CMD-WRKACTJOB	RUN
	QSERVER	QSYS	SBS	0		DEQW
	QPWFSERVSD	QUSER	BCH	0		SELW
	QPWFSERVS0	QUSER	PJ	0		DEQW

(Calculate the average of the CPU utilization for all job types "INT" for the desired time interval for interactive CPU utilization - "P" in the formula shown below.)

Another method is to run the Performance Monitor (STRPFRMON, ENDPFRMON) during selected time periods and review the first page of the Performance Tools/400 licensed program Component Report. The following is an example of this section of the report:

Component Report Component Interval Activity Data collected 190396 at 1030

Member . . . : Q960791030 Model/Serial . : 310-2043/10-0751D Main St... Library. . : PFR System name. . : TEST01 Version/Re..

ITV	Tns/hr	Rsp/Tns	CPU %	CPU%	CPU %	Disk I/O	Disk I/O
End		_	Total	Inter	Batch	per sec	per sec
						Sync	Async
10:36	6,164	0.8	85.2	32.2	46.3	102.9	39
10:41	7,404	0.9	91.3	45.2	39.5	103.3	33.9
10:46	5,466	0.7	97.6	38.8	51	96.6	33.2
10:51	5,622	1.2	97.9	35.6	57.4	86.6	49
10:56	4,527	0.8	97.9	16.5	77.4	64.2	40.7
:							
11:51	5,068	1.8	99.9	74.2	25.7	56.5	19.9
11:56	5,991	2.4	99.9	46.8	45.5	65.5	32.6

Itv End-----Interval end time (hour and minute)

Tns/hr-----Number of interactive transactions per hour

Rsp/Tns-----Average interactive transaction response time

(Calculate the average of the CPU utilization under the "Inter" heading for the desired time interval for interactive CPU utilization - "P" in the formula shown below.)

It is possible to have interactive jobs that do not show up with type "INT" or in the Performance Monitor Component Report. An example is a job that is submitted as a batch job that acquires a work station. These jobs should be included in the interactive CPU utilization count.

Most systems have peak workload environments. Care must be taken to insure that peaks can be contained in server model environments. Some environments could have peak workloads that exceed the interactive capacity of a server model or could cause unacceptable response times and throughput.

In the following equations, let the interactive CPU utilization of the existing traditional system be represented by percent P. A server model that should then produce the same response time and throughput would have a CPW of:

Server Interactive CPW = 3* P * Traditional CPWor for Custom Models use:* P * Traditional CPW (when P < 85%)</td>or* P * Traditional CPW (when P >= 85%)

Use the 1.5 factor to ensure the custom server is sized less than 85% CPU utilization.

These equations provide the server interactive CPU cycles required to keep the interactive utilization at or below the knee of the curve, with the current interactive workload. The equations given at the end of the Server and Custom Server Model Behavior section can be used to determine the effective interactive

utilization above the knee of the curve. The interactive workload below the knee of the curve represents one third of the total possible interactive workload, for non-custom models. The equation shown in this section will migrate a traditional system to a server system and keep the interactive workload at or below the knee of the curve, that is, using less than two thirds of the total possible interactive workload. In some environments these equations will be too conservative. A value of 1.2, rather than 1.5 would be less conservative. The equations presented in the **Interactive Utilization** section should be used by those customers who understand how server models work above the knee of the curve and the ramifications of the V4R1 enhancement.

These equations are for migration of "existing workload" situations only. Installation workload projections for "initial installation" of new custom servers are generally sized by the business partner for 50 - 60% CPW workloads and no "formula increase" would be needed.

For example, assume a model 510-2143 with a single V3R6 CPW rating of 66.7 and assume the Performance Tools/400 report lists interactive work CPU utilization as 21%. Using the previous formula, the server model must have an interactive CPW rating of at least 42 to maintain the same performance as the 510-2143.

Server interactive CPW = 3 * P * Traditional CPW = 3 * .21 * 66.7 = 42

A server model with an interactive CPW rating of at least 42 could approximate the same interactive work of the 510-2143, and still leave system capacity available for client/server activity. An S20-2165 is the first AS/400e series with an acceptable CPW rating (49.7).

Note that interactive and client/server CPWs are not additive. Interactive workloads which exceed (even briefly) the knee of the curve will consume a disproportionate share of the processing power and may result in insufficient system capacity for client/server activity and/or a significant increase in interactive response times.

2.12 Migration from Server Models

The section describes a recommended methodology for migrating from a server to a traditional model.

First determine the interactive CPU utilization for the server model. The second step is to determine the batch (client/server) CPU utilization for the server model. The previous section ("Migration from Traditional Models") describes how the Work with Active Jobs (WRKACTJOB) command or the Performance Monitor (STRPFRMON, ENDPFRMON) may be used to gather this information. The last step is to get the Maximum Client/Server CPW rating for the server.

Now in the following equations, let:

I = Interactive CPU UtilizationB = Batch CPU Utilization

CPWcs = Maximum Client/Server CPW rating

A traditional model that should produce the same response time and throughput would have a CPW of: Traditional CPW = CPWcs *(I + B) / .70

Note: In the above formula the division by 70 percent (.70) is done as a guideline to keep the system's CPU utilization at 70 percent, or less.

For example, assume a model 170-2160 with a V4R2 Maximum Client/Server CPW rating of 114, and assume the Performance Tools/400 report lists interactive work CPU utilization as 10% and batch CPU utilization at 50%. Using the previous formula, the traditional model should have a CPW rating of at least 97.7 to maintain the same performance as the 170-2160, this corresponds to an AS/400e 620-2180 system.

Traditional CPW = CPWcs *(I + B) / .70= 114 * (.10 + .50) / .70= 97.7

This formula should ensure that this system will give similar performance, however; each situation is unique and should be evaluated with an understanding of what the performance goals are. For example, if longer batch execution times are acceptable then a system with a lower CPW rating may be sufficient.

2.13 AS/400e Dedicated Server for Domino Performance Behavior

This section describes the performance behavior of three new processor features for the AS/400e server 170 which are marketed as the AS/400e Dedicated Server for Domino (generally available September 24, 1999). The AS/400e Dedicated Server for Domino (DSD) exloits AS/400 system resources to deliver improved price/performance for Lotus Domino workloads. For additional description of what kinds of workloads are considered Domino versus non-Domino, please refer to section "AS/400e Dedicated Server for Domino" in Chapter 11, "Domino for the AS/400".

Announced DSD capabilities:

•	170-2407	Simple Mail Users $= 1,300$	Processor CPW = 30	Interactive $CPW = 10$
•	170-2408	Simple Mail Users $= 2,300$	Processor CPW = 60	Interactive $CPW = 15$
•	170-2409	Simple Mail Users $= 4,300$	Processor $CPW = 120$	Interactive $CPW = 20$

Model 170 scalability (measured as Simple Mail Users) has not increased with the introduction of the three new processor features that are based on existing Northstar technology. However, the new processor features available for the AS/400e Dedicated Server for Domino can support a much higher number of Simple Mail Users when compared to standard processors offered on the model 170 with similar CPW ratings. This re-balancing of workload processing capabilities on the DSD results in attractive price/performance levels for customers seeking a reliable and secure server that focuses on multiple Domino workloads.

The following charts and examples describe the behavior of Domino and non-Domino workloads on the DSD processor features.

DSD Interactive CPW capacity

The DSD processor features will support reasonable interactive processing required to perform system administration functions. The DSD processor features are not intended to run sustained interactive workloads and will be subject to the interactive CPW capabilities described above. In general, DSD will support interactive processing up to approximately 2-3% CPU utilization before hitting the "knee of the curve". Beyond the knee of the curve, the system will reach 100% CPU when the maximum amount of interactive CPU processing (approximately 3-5%) is reached. Please refer to **Figure 2.4. Model 7xx Utilization Example** in Chapter 2, **Section 2.4. Performance Highlights of New Model 7xx Servers** for a more thorough description of interactive behavior on server models. The discussion in Section 2.4 applies to DSD. Earlier sections in Chapter 2 discuss other server model behaviors which also apply to DSD including:

- knee of the curve
- interactive and client/server CPW values not being additive (interactive processing is part of non-Domino client/server processing)
- server dynamic tuning

DSD Processor (client/server) CPW capacity

A limited amount of non-Domino client/server processing is available on a DSD. This capacity is intended to support a limited amount of system resource activity (Integrated File System, communications, storage management, back-up and recovery, etc.) and Domino application integration functions (DB2 Universal Database access, external program calls, Java applications, etc.) in support of the Domino application operating on the server. As such, this capacity does not represent a guaranteed level of capacity to perform non-Domino work.

The non-Domino workload on a DSD should be managed to not exceed the Processor CPW rating (approximately 10-15% of the total CPU capacity). Beyond the Processor CPW rating the CPU will approach 100% and eventually saturate. The point at which this occurs varies by processor feature code.

For normal operations on a DSD, system overhead (seen in CFINT) will generally be at small, nominal values. However, for cases where the Domino workload is unusually small and non-Domino client/server processing is present, larger amounts of system overhead may occur. *To ensure the most efficient use of CPU resources on a DSD, a good rule of thumb is to manage the ratio of Domino processing to non-Domino client/server CPU processing such that it remains above 3 to 1*. Below this ratio, a disproportionately large amount of client/server processing versus Domino processing will result in increased system overhead.

Performance Tip: Use job priority settings to prioritize Domino processing above non-Domino client/server processing. This will allow the Domino workload to compete more effectively for system resources and lessen the amount of system overhead as discussed in the examples below. This is particularly important at high levels of CPU utilization.

The following figures describe DSD behavior for the 170-2409 processor feature for four different workload environments and assess whether they are appropriate for DSD. (Note: The 2407 and 2408 processor features will allow a slightly higher amount of non-Domino client/server processing than what is shown for 2409). The examples are arranged as follows:

1. **DSD appropriate workload balance** (Domino to non-Domino client/server ratio >3:1) A. Domino workload remains constant, non-Domino client/server workload varies

- B. Domino and non-Domino client/server workloads increase proportionately
- DSD marginal workload balance (Domino to non-Domino client/server ratio slightly less than 3:1)
 A. Domino workload remains constant, non-Domino client/server workload varies
 B. Domino and non-Domino client/server workloads increase proportionately
- DSD inappropriate workload balance (Domino to non-Domino client/server ratio <3:1)
 A. Domino workload remains constant, non-Domino client/server workload varies
 B. Domino and non-Domino client/server workloads increase proportionately
- 4. DSD "no Domino" workload (an example of system administration when Domino is not active)

Please note that in all of the following examples most of the system overhead that is shown is processing that can be reclaimed by increasing the amount of Domino processing on the system.

1. DSD appropriate workload balance (greater than 3 to 1 ratio of Domino to non-Domino client/server)

In Figures 2.7 and 2.8 below, DSD system resources are being used as efficiently as possible because the ratio of Domino processing to non-Domino client/server processing is greater than 3 to 1.







B. Domino and non-Domino processing increase proportionately



Figure 2.8

In Figures 2.7 and 2.8 above, when operating with higher than recommended amounts of non-Domino client/server processing (i.e. operating near the right edge of the graph), it is particularly important to have the priority of the Domino processing higher than that of the non-Domino processing. In such a heavily loading case, the higher relative priority of the Domino processing will help ensure that it continues to receive an appropriate amount of the processing resources, with minimal interference from the non-Domino client/server processing or system overhead.

2. DSD marginal workload balance (near a 3 to1 ratio of Domino to non-Domino client/server)



A. Domino processing is constant at 45% CPU, non-Domino client/server processing varies



In Figure 2.9 above, near a point of 15% non-Domino client/server processing the ratio drops below 3 to 1 and system overhead becomes measurable as the non-Domino workload increases. Note that for both examples in Figure 2.9 above and Figure 2.10 below that if the Domino workload were to increase and were a higher priority than the non-Domino client/server workload, the Domino workload would reclaim the resources from the non-Domino workload and its associated system overhead.

B. Domino and non-Domino client/server processing increase proportionately



Figure 2.10

3. DSD inappropriate workload (below a 3 to 1 ratio of Domino to non-Domino client/server)

In Figures 2.11 and 2.12 below, note that system overhead is present at a low level of non-Domino client/server processing because far too little Domino processing is occurring to effectively utilize the server. If the Domino workload were to increase and were a higher priority than the non-Domino client/server workload in this example, the Domino workload would reclaim the resources from the non-Domino workload and its associated system overhead.

A. Domino processing is constant at 20% CPU, non-Domino client/server processing varies





B. Domino and non-Domino client/server processing increase proportionately



Figure 2.12

V4R4 Performance Capabilities Reference 8 Copyright IBM Corp. 1999 4. DSD "no Domino" workload (example of system administration when Domino is not active)

This example depicts an environment where no Domino workload is present at all, such as would possibly be the case during system administration activities such as back-up. Adequate system resources are available to perform reasonable administration activities such as system back-up, even though system overhead is present. Please note that even though a large amount of system overhead may be present in such cases, the server is providing the full rated non-Domino client/server Processor CPW for such operations.



Figure 2.13
Chapter 3. Batch Performance

In a commercial environment, batch workloads tend to be I/O intensive rather than CPU intensive. The factors that affect batch throughput for a given batch application include the following:

- Memory (Pool size)
- CPU (processor speed)
- DASD (number and type)
- System tuning parameters

Batch Workload Description

The Batch Commercial Mix is a synthetic batch workload designed to represent multiple types of batch processing often associated with commercial data processing. The different variations allow testing of sequential vs random file access, changing the read to write ratio, generating "hot spots" in the data and running with expert cache on or off. It can also represent some jobs that run concurrently with interactive work where the work is submitted to batch because of a requirement for a large amount of disk I/O.

3.1 Effect of CPU speed on Batch

The capacity available from the CPU affects the run time of batch applications. More capacity can be provided by either a CPU with a higher CPW value, or by having other contending applications on the same system consuming less CPU.

Conclusions/Recommendations

- For CPU-intensive batch applications, run time scales inversely with Relative Performance Rating (CPWs). This assumes that the number synchronous disk I/Os are only a small factor.
- For I/O-intensive batch applications, run time may not decrease with a faster CPU. This is because I/O subsystem time would make up the majority of the total run time.
- It is recommended that capacity planning for batch be done with tools that are available for AS/400. For example, BEST/1 for OS/400 (part of the Licensed Program Product, AS/400 Performance Tools) can be used for modeling batch growth and throughput. BATCH400 (an IBM internal tool) can be used for estimating batch run-time.

3.2 Effect of DASD type on Batch

For batch applications that are I/O-intensive, the overall batch performance is very dependent on the speed of the I/O subsystem. Depending on the application characteristics, batch performance (run time) will be improved by having DASD that has:

- faster average service times
- read ahead buffers
- write caches

Additional information on DASD devices in a batch environment can be found in Chapter 14, "DASD Performance"0.

3.3 Tuning Parameters for Batch

There are several system parameters that affect batch performance. The magnitude of the effect for each of them depends on the specific application and overall system characteristics. Some general information is provided here.

• Expert Cache

Expert Cache did not have a significant effect on the Commercial Mix batch workload. Expert Cache does not start to provide improvement unless the following are true for a given workload. These include:

- the application that is running is disk intensive, and disk I/O's are limiting the throughput.
- the processor is under-utilized, at less than 60%.
- the system must have sufficient main storage.

For Expert Cache to operate effectively, there must be spare CPU, so that when the average disk access time is reduced by caching in main storage, the CPU can process more work. In the Commercial Mix benchmark, the CPU was the limiting factor.

However, specific batch environments that are DASD I/O intensive, and process data sequentially may realize significant performance gains by taking advantage of larger memory sizes available on the RISC models, particularly at the high-end. Even though in general applications require more main storage on the RISC models, batch applications that process data sequentially may only require slightly more main storage on RISC. Therefore, with larger memory sizes in conjunction with using Expert Cache, these applications may achieve significant performance gains by decreasing the number of DASD I/O operations.

• Job Priority

Batch jobs can be given a priority value that will affect how much CPU processing time the job will get. For a system with high CPU utilization and a batch job with a low job priority, the batch throughput may be severely limited. Likewise, if the batch job has a high priority, the batch throughput may be high at the expense of interactive job performance.

• Dynamic Priority Scheduling

See 19.2, "Dynamic Priority Scheduling" Ofor details.

• Application Techniques

The batch application can also be tuned for optimized performance. Some suggestions include:

- Breaking the application into pieces and having multiple batch threads (jobs) operate concurrently. Since batch jobs are typically serialized by I/O, this will decrease the overall required batch window requirements.
- Reduce the number of opens/closes, I/Os, etc. where possible.

- If you have a considerable amount of main storage available, consider using the Set Object Access (SETOBJACC) command. This command pre-loads the complete database file, database index, or program into the assigned main storage pool if sufficient storage is available. The objective is to improve performance by eliminating disk I/O operations.
- If communications lines are involved in the batch application, try to limit the number of communications I/Os by doing fewer (and perhaps larger) larger application sends and receives. Consider blocking data in the application. Try to place the application on the same system as the frequently accessed data.

3.4 V4R4 comments

We observed an increase in the CPU requirements for traditional (RPG and COBOL end-of-day processing) batch workloads of 5-8%. Except for environments where the system is nearing the need for an upgrade or environments where a particular job must finish prior to other jobs starting, we do not expect this to have a major effect on the overall batch window.

Chapter 4. DB2 UDB for AS/400 Performance

This chapter contains performance information on items that are important to achieving a good overall level of performance for DB2 UDB for AS/400 (previously DB2 for AS/400) environments. The information presented here concentrates on performance for applications run locally on DB2 UDB for AS/400, although much of the information can also be used to ensure better levels of performance for applications using remote access to an AS/400.

The first section in this chapter provides information on what has changed in DB2 UDB for AS/400 in V4R4. The second section concentrates on DB2 enhancements made in previous Version 4 releases (V4R3, V4R2 and V4R1). The last section contains articles that were written in past (Version 3) releases about DB2 for AS/400 performance characteristics.

0

4.1 V4R4 Enhancements for DB2 UDB for AS/400

In V4R4, DB2 for AS/400 has now been renamed to DB2 Universal Database for AS/400 (DB2 UDB for AS/400). Although many of the enhancements in V4R4 for DB2 UDB for AS/400 are functional in nature, there are also significant changes that will help improve query performance in certain key areas. This section will describe what has changed from a performance perspective in V4R4. If you want to know more about overall changes for DB2 UDB for AS/400, refer to the items described later in this section under the title "Sources of Additional Information".

Performance Improvements in V4R4

Following is a description of the changes that were made to help improve query performance in V4R4. Enhancements were made in this release to help improve performance for both business intelligence (BI) queries as well as for those queries likely to be found in everyday batch and OLTP applications. Although performance measurement data is not available for these enhancements, these changes will in many cases result in noticeable improvements for queries that are able to take advantage of them.

1) General improvement in query optimization and runtime system programs

Many of the query optimization and runtime programs have been changed to use a new IBM internal programming language to take advantage of better compiler optimization. This has resulted in noticable performance improvements in some areas, particularly in the amount of CPU used during query optimizations.

2) Reduction in memory used by queries

Changes have been made to internal space allocation algorithms and to internal query structures that will help reduce the memory footprint for most queries. This will help reduce overall memory consumption, especially in environments running with large numbers of reusable ODPs.

3) Improved run times for some temporary sorts

On the AS/400, queries using temporary sorts often show a relatively high cost at open time, even when the open is for a reusable ODP. A significant portion of this cost is due to initializations required for the sort, and in cases where the sort only involves a few records, the initialization cost accounted for a large portion of the open. In V4R4, temporary sorts involving a small number of rows will now use a different sort algorithm that will significantly reduce this cost at open time. For environments that involve a significant number of reusable ODPs with small temporary sorts, this enhancement could result in noticable improvements in overall run time.

4) Increased use of hash join

The use of a hash join will in many cases be more efficient than a nested loop join. However, prior to V4R4, hash joins were not allowed for queries running with commitment control levels of *CHG or *CS or for queries with subqueries where the entire query was implemented as a composite join. In V4R4, these restrictions have been removed, which may result in improved performance for those queries which previously ran under these limitations but could have benefited from using a hash join.

5) Improved performance for MIN and MAX functions

In V4R4, queries that involve use of the grouping functions MIN or MAX for a given column will in some cases realize a performance improvement provided that the grouped column is retrieved using an index. To find out more about what types of queries will benefit from this change, refer to the DB2 UDB for AS/400 SQL Programming manual (SC41-5611) for V4R4.

6) Improvements for partial outer and exception joins

In V4R4, the optimizer can now implement partial outer (PO) joins and exception (EX) joins using using a join method known as the key positioning access method. This join method existed prior to V4R4 and was used in other join types, but were not allowed in PO or EX joins. Use of the key positioning access method in these join types and in queries with mixtures of these joins can in many cases result in significant performance improvements.

7) Better use of EVIs and dynamic bitmaps

In V4R4, changes have been made to allow improved use of EVIs and dynamic bitmaps. For example, multiple EVIs built over both the fact and dimension tables can be used to generate a list of all possible values to be selected from the fact table. This list is then used to generate a bitmap that can be used in scanning the fact table, which can result in noticable performance gains. Performance of some join queries, in particular star schema joins, may also show significant improvements from these changes.

8) Improved internal handling of SQL opens and cursors

Changes have been made in V4R4 to enable parameter marker conversion to occur in more cases, which will help reduce the number of full opens that occur in some dynamic or extended dynamic environments. Also, improvements have been made to the algorithms used to find an existing open cursor (reusable ODP) within a job. For jobs with a large number of reusable ODPs, this improvement can be noticeable.

9) Improved interface for debugging query performance

A new INI file interface is available in V4R4 that provides users with the ability to dynamically modify or override the environment in which queries are executed. This support is shipped with the V4R4 base release. This interface also allows Rochester lab developers the ability to control and debug queries for performance without having to install additional tools on the system. For more information on this interface and how to use it, refer to the DB2 UDB for AS/400 SQL Programming manual (SC41-5611) for V4R4.

Sources of Additional Information

1) For further information on other DB2 UDB for AS/400 enhancements in V4R4, see the Internet page at the following url:

http://www.as400.ibm.com/db2

This page contains references and urls that can be followed to obtain information such as recent announcements, articles and white papers, technical information and tips, and further enhancements and performance improvements for DB2 UDB for AS/400.

In addition, the following urls point to specific articles that provide additional functional and performance information about V4R4 changes for DB2 UDB for AS/400:

http://www.as400.ibm.com/developer/comm/newsletter.html

http://www.news400.com/NWN/StoryBuild.cfm?ID=436

2) A course for developing DB2 UDB for AS/400 SQL and query performance tuning skills will be made available shortly. The following url will be updated with information on this course as it becomes available:

http://www.as400.ibm.com/db2/db2educ_m.htm

3) Some of the above enhancements as well as other potential query performance improvements have also been made available in what are called database fixpack PTF packages for V4R3 and V4R2. To find out more about these fixpacks, refer to the item on these fixpacks at the following url:

http://www.as400.ibm.com/db2/db2tch_m.htm

In addition to fixpack information, this url points to other items such as technical overview articles, publications and redbooks and other enhancements, both functional and performance related.

4) The AS/400 Teraplex Center plays a significant role in verifying the benefits of new technologies for data warehousing operations. In many cases their testing applies to other general applications of these

technologies as well. For a variety of recent test results and additional information, on many of the new see the Teraplex Center's Internet home page at url:

http://www.as400.ibm.com/developer/bi/teraplex

4) The AS/400 Systems Performance area maintains an IBM intranet home page (available to IBM personnel and others who have access to the IBM intranet). For a variety of white papers and additional information on many of the technologies described in this chapter, see the Performance section at url:

http://ca-web.rchland.ibm.com/perform/perfmenu.htm

5) For information and guidance about database performance tuning and query performance tuning, refer to the DB2 UDB for AS/400 Database Programming manual (SC41-5701) (appendix on query performance) and the DB2 UDB for AS/400 SQL Programming manual (SC41-5611) (chapter on data management and query optimizer).

4.2 Previous Version 4 Enhancements for DB2 for AS/400

Encoded Vector Indices (EVIs)

In V4R3 a new type of permanent index, the Encoded Vector Index (EVI), can be created through SQL. EVIs cannot be used to order records, but in many cases, they can improve query preformance. An EVI has several advantages over a traditional binary radix tree index.

- The query optimizer can scan EVIs and automatically build dynamic (on-the-fly) bitmaps much more quickly than traditional indexes. For more information on dynamic bitmaps, see the description in their section below.
- EVIs can be built much faster and are much smaller than traditional indexes. Smaller indexes require less DASD space and also less main storage when the query is run.
- EVIs automatically maintain exact statistics about the distribution of key values, whereas traditional indexes only maintain estimated statistics. These EVI statistics are not only more accurate, but also can be accessed more quickly by the query optimizer.

EVIs are used by the AS/400 query optimizer with dynamic bitmaps and are particularly useful for advanced query processing. EVIs will have the biggest impact on the complex query workloads found in business intelligence solutions and adhoc query environments. Such queries often involve selecting a limited number of rows based on the key value being among a set of specific values (eg a set of state names).

When an EVI is created and maintained, a symbol table records each distinct key value and also a corresponding unique binary value (the binary value will be 1, 2, or 4 bytes long, depending on the number of distinct key values) that is used in the main part of the EVI, the vector (array). The subscript of each vector (array) element represents the relative record number of a database table row. The vector has an

entry for each row. The entry in each element of the vector contains the unique binary value corresponding to the key value found in the database table row.

The following is an example of how to create an EVI with the SQL CREATE INDEX statement:

CREATE ENCODED VECTOR INDEX StateIx ON CUSTOMERS (CustState)

Parallel Data Loader

The data loader is a new function in V4R3 (also available via PTF on V4R2 and V4R1) that makes loading AS/400 database tables from external data much simpler and faster. The data loader can import fixed-format, delimited, and byte-stream files. A new CL command, Copy From Import File (CPYFRMIMPF), is provided to simplify the process.

After installing and activating the DB2 for AS/400 SMP licensed feature on a multiprocessor AS/400, parallel processing increases the speed of the data loader by approximately ten times over non-parallel methods. With this feature active, DB2 for AS/400 is able to use multi-tasking, rather than just a single task, to load an import file.

The following PTFs are required to provide these functions on earlier releases:

- V4R1M0 PTFs: SF47138 and SF47177 for OS/400
- V4R2M0 PTFs: SF46911 and SF46976 for OS/400

These are available individually and may be available in a cumulative\pard PTF package. **Parallel Index Maintenance**

Parallel index maintenance, supported in V4R3, can be useful to those that have many logical files and indexes defined over a single database file. Every time that a row is inserted, changed or deleted into a database table, all of the indexes and logical files defined over that base table have to be maintained to reflect the latest data change. The parallel index maintenance enhancement allows DB2 for AS/400 to maintain multiple indexes in parallel instead of one at a time as done in previous releases. Note that DB2 for AS/400 will utilize parallel index maintenance only when blocked insert operations are performed on the base database table, there exist at least eight indexes over the table, and the DB2 for AS/400 SMP (Symmetric Multiprocessing) licensed feature is installed and activated. Parallel index maintenance thus allows DB2 for AS/400 to reduce the amount of time it takes to maintain indexes when you are adding lots of new rows to a database table. The data loader and copy file (CPYF) utilities also will benefit from this feature since they utilize blocked inserts.

Dynamic Bitmaps

Dynamic bitmaps, introduced in V4R2, can improve the performance of certain query operations. A dynamic (on-the-fly) bitmap is a temporary binary structure built against a permanent index. The AS/400 database query optimizer automatically builds dynamic bitmaps when the optimizer determines that dynamic bitmaps will speed up response time. Use of dynamic bitmaps allows the system to perform skip-sequential DASD operations and reduces the need for full-table-scans, thereby reducing database I/O

operations and speeding up completion of the affected queries. In addition, DB2 for AS/400's dynamic bitmap support allows the use of multiple indexes against any particular table in the query (previously limited to at most one index per table). These multiple bitmaps resulting from the use of more than one index are combined into a composite results bitmap using boolean logic.

When the optimizer builds a dynamic bitmap, it sets a bit for every index entry that meets the selection criteria. It can combine (AND, OR) multiple bitmaps into a composite results bitmap. The optimizer uses this final bitmap to retrieve only those records whose bits are set.

Note that single-column indexes can often maximize flexibility for the database administrator. Dynamic bitmaps allow the optimizer to use several of these single-column indexes at once and combine their dynamic bitmaps, using boolean logic, into one bitmap. In this way, ad hoc users of large databases may be able to realize large performance gains without significant impact to the system and with a minimal set of indexes.

System-Wide SQL Statement Cache

The system-wide SQL statement cache, introduced in V4R2, can improve the performance of programs using dynamic SQL. The system automatically caches dynamic SQL statements. No user action is required to activate or administer this function.

When a dynamic SQL statement that has previously executed is later reexecuted, if the statement is still available in the cache, DB2 for AS/400 can retrieve (rather than construct) key information associated with the cached statement, and thereby reduce the processing resource and the time required to execute the statement again.

Remote Journal Function

Introduced in V4R2, the remote journal function allows replication of journal entries from a local (source) AS/400 to a remote (target) AS/400 by establishing journals and journal receivers on the target system that are associated with specific journals and journal receivers on the source system. Some of the benefits of using remote journal include:

- Allows customers to replace current programming methods of capturing and transmitting journal entries between systems with more efficient system programming methods. This can result in lower CPU consumption and increased throughput on the source system.
- Can significantly reduce the amount of time and effort required by customers to reconcile their source and target databases after a system failure. If the synchronous delivery mode of remote journal is used (where journal entries are guaranteed to be deposited on the target system prior to control being returned to the user application), then there will be no journal entries lost. If asynchronous delivery mode is used, there may be some journal entries lost, but the number of entries lost will most likely be fewer than if customer programming methods were used due to the reduced system overhead of remote journal.
- Journal receiver save operations can be off-loaded from the source system to the target system, thus further reducing resource and consumption on the source system.

Hot backup, data replication and high availability applications are good examples of applications which can benefit from using remote journal. Customers who use related or similar software solutions from other vendors should contact those vendors for more information.

Test Environment

As mentioned above, remote journal can be run in either synchronous delivery mode or asynchronous delivery mode. In synchronous mode, the source system must wait for a confirmation message from the target system that the journal entries have been received. In asynchronous mode, the source system runs without having to wait for remote journal to finish. In both sync and async modes, remote journal can be in one of two activation modes. If the source journal receivers contain entries that have not been replicated at the time remote journal is started, then remote journal will run in "catch-up" mode to transfer these entries to the remote system. Once remote journal catches up with the source system journal entries, it then runs in "continuous" mode.

Tests were done in Rochester to evaluate the following aspects of remote journal performance:

- Comparing elapsed times of running catch-up mode using TCP/IP, APPC and Opticonnect for OS/400
- Overall impact of running remote journal in continuous mode in an interactive transaction processing environment using TCP/IP, APPC and Opticonnect for OS/400

For all tests, the source system used was a model 530-2151 with 2.5 GB of memory and 73 GB of DASD (24 arms totalling 65 GB in the system ASP, 4 arms totalling 8 GB in a user ASP). The target system was a model 510-2143 with 1.0 GB of memory and 33 GB of DASD (16 arms totalling 25 GB in the system ASP, 4 arms totalling 8 GB in a separate user ASP). All DASD arms on both systems were unprotected (RAID or mirroring was not used). Both systems were installed with V4R2, and on both systems, the journal receivers were located on the user ASP.

For both the TCP/IP and APPC tests, a 16Mbps token ring was used (model 2629 with 6149 cards on both systems). For Opticonnect, the bus model used was a 2682, with the source system using a 2685 card and the target system using a 2688 card.

Note that the results shown here are for specific environments and configurations running a controlled and repeatable series of tests. The actual results you obtain in your environment may vary from what is presented here, although the conclusions and recommendations made here will be applicable to most customer applications using remote journal.

Remote Journal Catch-Up Mode

For the catch-up mode tests, approximately 3.1 GB of journal entries was transferred to the remote system when remote journal was started. Memory pools were allocated to ensure this resource was not constrained. There was no other activity on either the source or target systems.

Using TCP/IP and APPC, this transfer took just over 26 minutes, while it took about 12 minutes using Opticonnect. These results are as expected given the much faster transfer rates of the Opticonnect bus versus the token ring connection. In all cases, the CPU utilization was low (15% in the Opticonnect case and less than 10% in the other measurements). No other system resources were constrained during these tests.

Remote Journal Continuous Mode

The base run for this test was done using an interactive transaction processing environment with 640 users running about 224,000 transactions per hour on simulated locally attached 5250 workstations. All files that were updated in this environment were journaled to a single local journal receiver on the source system's user ASP. Commitment control was not used. This environment produces about 3.2 million journal entries per hour with a CPU utilization of about 72% on the source system. Memory and DASD are not constrained in this environment. The response time for this environment is based on one transaction type that accounts for 45% of the total throughput and produces about 35 journal entries per transaction. In this environment, the average response time for this transaction was 0.15 seconds.

For the remote journal runs, all journal entries produced on the source system were replicated on the target system. Both asynchronous mode and synchronous mode were tested on each of the three different communications protocols. In the remote journal runs, the target journal receiver was located on the user ASP on the target system, and again memory and DASD were not constrained. There was no other activity on the target system.

When remote journal with asynchronous mode was added to the base environment, the results show that for all three protocols, the impact to response time was minimal (about 0.05 seconds) with a corresponding drop in throughput of less than 1%. The increase in CPU utilization on the source system ranged from 2-7%, and the token ring and bus utilizations were very reasonable (less than 15% in all cases). The target system in all three environments showed a CPU utilization of 2-3% with low levels of disk arm utilization.

When remote journal with synchronous mode was added to the base environment, the Opticonnect run showed the least impact to response time (increase of about 0.15 seconds) with a corresponding drop in throughput of about 1.5%, while TCP/IP and and APPC showed an increase of about 0.3 seconds with a drop in throughput of about 3%. In all three environments, the CPU overhead per transaction from adding remote journal was 7-8% on the source system. The token ring utilization in the APPC and TCP runs was about 15-20%, and the bus utilization was also low in the Opticonnect run. The CPU overhead on the target system was about 5-7%, and DASD utilizations on the target system user ASP were again low as they were in the async mode runs.

Conclusions and Recommendations

• For most interactive environments, adding remote journal should not result in significant degradations to response time or resource utilization on either the source or target system. Although you can expect a more noticeable increase in response times when using synchronous mode, the tests above show that the amount of the increase should still be reasonable. In sync mode, the amount of increase will be relative to the number of journal entries produced by an average transaction.

One other item to consider is that the interactive transactions described here spend over 70% of their time doing database activity, with 20% or so spent in application code. In many customer environments, the average transaction spends a much higher percentage of time in application code and other areas and a much lower percentage in database activity. Overall, this means that for many customer applications, the overall impact from adding remote journal may be less than it was in the tests described here.

There are several factors to consider when deciding what sort of remote journal setup you will choose:

- 1. Whether or not you can afford loss of some journal entries in the event of a system failure. Synchronous mode will guarantee no journal entries are lost, while asynchronous mode may lose some entries (although the number will be relatively low compared to the total number of journal entries being replicated).
- 2. In general, Opticonnect offers noticeable advantages over APPC or TCP/IP alternatives via token ring, including less of an impact to response times (particularly when using sync mode) over a broader range of throughput and also higher overall capacity. For customers who already have Opticonnect installed and have additional capacity available, this would be a logical choice for implementing remote journal. However, users who do not currently have Opticonnect need to balance the added cost of this product versus what can be accomplished over a token ring using APPC or TCP/IP. For example, if the token ring has enough capacity for your level of remote journal and you either use async mode or can afford the response time increase with sync mode, then using APPC or TCP/IP with the token ring would also be a logical choice.
- 3. Another item to consider for remote journal is DASD performance, particularly when using Opticonnect. Since Opticonnect allows a much higher transfer rate and capacity without bottlenecking the bus, the performance of the DASD arms where the journal receiver is located can become a problem in very heavy journaling environments, particularly if the arms are RAID protected. In heavy remote journal environments, Opticonnect can achieve transfer rates up to 1 GB per second when the journal receivers on both systems are located on unprotected DASD arms (no RAID or mirroring). If RAID protection is used in the same environment, the DASD arms become the bottleneck. If protection is required, using mirrored arms for the journal receiver DASD will provide significantly better performance than RAID protection. However, unprotected DASD arms will allow the highest overall transfer rates and best overall performance when using Opticonnect for remote journal. Whatever implementation is used, DASD arm performance should still be monitored using the performance monitor and/or other tools such as the WRKDSKSTS command.

In most cases involving high transfer rates via TCP/IP or APPC, the communications line or IOP will tend to become the bottleneck prior to the journal receiver DASD arms becoming overutilized. However, it is still a good idea to monitor the performance of these arms to ensure the best overall results.

- 4. In catch-up mode, performance will be significantly better using Opticonnect than with TCP or APPC due to the higher transfer rates and overall higher capacity of the Opticonnect bus. This should also be considered when determining what medium to use in your overall remote journal implementation.
- 5. The measurements done here show the impact of adding remote journal to an application that is doing local journaling without commitment control. Similar results can be expected when adding remote journal to customer applications that use commitment control.
- If remote journal will be used during execution of a batch job, it is recommended that asynchronous mode be used. The reason for this is that in the event of a system failure, most customers choose to start the batch job over from the beginning instead of restarting the batch job partway through after recovering to that point using the remote journal. Because of this, the additional overhead of using

synchronous mode will result in longer elapsed times with no additional benefits.

• As noted several times in this section, it will be important to monitor your resource utilizations on the source and target systems both prior to and after implementing remote journal. Key resources to consider include CPU utilization, line utilization, IOP utilizations, and DASD arm performance, as well as overall response times. These items can be monitored effectively using data collected from the performance monitor tool.

Additional Sources of Information

Customers interested in using the remote journal function should read the chapter on remote journal in the V4R2 version of the *OS/400 Backup and Recovery Guide*, SC41-5304. This chapter contains additional information on functional and performance aspects of remote journal.

Parallel Index Build

Parallel index build, introduced in V4R1, has several important uses. This function can be used when the DB2 for AS/400 SMP feature is installed and made active. Following is a list of its uses.

- Speeding up the process of building a temporary index, when the query optimizer cannot locate an appropriate index and determines to build a temporary index to execute a particular query. This function is most beneficial in data warehousing environments housing large database files.
- Speeding up the process of building permanent indexes, particularly when built over large database files. Tests at the AS/400 Teraplex Center show that parallel index build, after the loading is complete, is the fastest way to load data and build an index, and is preferred to building indexes while data is being loaded.

V4R3 note: In V4R3, the new encoded vector indexes, as well as the traditional binary radix tree indexes, are supported by the parallel index build function.

4.3 Version 3 DB2 for AS/400 Performance Information

Although much of the information in this section is still applicable to DB2 for AS/400, there may be portions of the articles that are no longer accurate due to changes made since the article was written. As of V4R3, the following general comments should be considered when using the information in this section.

• Prior to V4R1, *MAX4GB was the default value when creating any new index (as discussed below in "Enhanced Index Support for DB2/400" on page 43). Starting in V4R1, the default has been changed to *MAX1TB. In addition, the potential performance concerns with *MAX1TB indexes have been alleviated such that *MAX1TB is now the recommended size for all indexes (hence the change to the default value). The only concern with using *MAX1TB is that indexes created with this value can be up to 15% larger in size than if *MAX4GB is used, so it may be best to monitor available DASD space when converting or creating these types of indexes.

- In addition to the SQL improvements discussed in sections "DB2/400 SQL and Query Information Improvements" on page 44 and "DB2/400 SQL" on page 46, there have been additional changes made in subsequent releases that in many cases will allow well-tuned SQL applications on the AS/400 to be competitive with these same applications run on other hardware platforms. Although there are too many changes to list or discuss here, it is recommended that customers currently running their SQL applications on V4R1 or prior releases consider moving to V4R2 or a later release to take advantage of these performance improvements. Users should be able to take advantage of these improvements with no changes to or recompiles of existing SQL applications.
- In the "DB2/400 SQL" on page 46 section, several of the additional sources of information listed may no longer be available under the same document numbers or titles. The titles listed in this section were accurate as of V3R6 and V3R7, but may have since been moved or deleted. It may be best to contact your IBM representative if you need to know what current sources of information are available concerning SQL performance.

DB2/400 Enhancements in V3R6

Enhanced Index Support for DB2/400

The index support was enhanced in V3R6 to support larger indexes and to reduce index seize contention. Now rather than a size limitation of 4 Gigabytes (GB), each index can be as large as One Terabyte (TB). This limit is related to the size of the index and not the number of entries.

To take advantage of the increased index size capabilities, a new parameter, ACCESS PATH SIZE, has been added to the Create Physical File (CRTPF), Create Logical File (CRTLF), and Change Physical File (CHGPF) commands. If you want to allow growth beyond 4 GB on a keyed physical file that already exists, you can use the Change Physical File (CHGPF) command specifying "*MAX1TB" for the ACCESS PATH SIZE parameter. If you want to change a logical file to the larger limit, you would delete the existing logical file and create a new logical file specifying "*MAX1TB" for the ACCESS PATH SIZE parameter. When creating new files, physical or logical, the default is *MAX4GB. You must specify *MAX1TB if it is required. It is recommended, though not required, that all access paths for a file be of the same type.

Note: The procedure described above will not work for files that have UNIQUE *YES specified. For these files you will have to create a second file with the ACCESS PATH SIZE specified at *MAX1TB and copy the first file to the second. The first file can then be deleted and the second file renamed.

In addition to the capability to create larger files, the algorithm used for seizing indexes has been enhanced. Indexes are now seized at the index page level. Therefore, for those workloads where there is a high-level of concurrency on a particular file or access path, the new algorithm will significantly reduce the contention resulting in significant performance improvements. This is particularly applicable to multi-processor systems and indexes with a high number of records being added. In order to take advantage of this change, the index must be created/changed as specified above.

Most customers will be best served by specifying the default of *MAX4GB for the ACCESS PATH SIZE parameter. This will in general provide better performance. Also if files will be moved to a prior release, the index may need to be rebuilt or the save of the file will not work if the index was created with *MAX1TB.

It is recommended that *MAX1TB be specified only where it is needed to allow for the larger file size or where there is high contention on access paths. If an index is changed to *MAX1TB and there is not high contention on the index, it may result in additional overhead. One measurement on a uni-processor system where this was the case resulted in a slight (less than 5%) slowdown. The amount of overhead would depend on, but not limited to, number of files, size of files, and access patterns.

How to determine when to switch for performance

To determine if the contention on your system is at a level that changing to a *MAX1TB access path will improve performance you will need to collect data using the Performance Monitor. This can be collected using the command STRPFRMON (Start Performance Monitor). When issuing this command, the interval value should be changed from 15 minutes to 5 minutes and the trace data collected should be changed from *NONE to *ALL. Data should be collected for at least 30 minutes and should be collected during "peak" activity. Please note that when the Performance Monitor is ended the trace data will automatically be dumped to DASD and performance could be degraded while this data is being transferred. This can be avoided by sepcifying Dump Trace *NO when starting the monitor and at a later time, when the system is not as busy, issuing the DMPTRCDTA command.

If the number of Seize Conflicts Per Second on the Component Report (created by issuing the PRTCPTRPT command or using the menu options from the PERFORM menu) is greater than 140 then you can in many cases benefit from changing to files with the ACCESS PATH SIZE parameter set to *MAX1TB. The method to determine which files should be changed is discussed in the next paragraph.

Once the trace data has been saved (either when the monitor ends if Dump Trace *YES was specified or after issuing the DMPTRCDTA command) a transaction report should be printed by either issuing the PRTTNSRPT command or using the menu options from the PERFORM menu. The SUMMARY OF SEIZE/LOCK CONFLICTS BY OBJECTS section of this report shows the number of seize conflicts and the number of lock conflicts by object. This list will show conflicts on both data spaces (ie. file data) and data space indexes (ie. access paths). If the majority of seize conflicts are on one or two data space indexes, those are the candidates for the larger access path size parameters. The files most likely to fall in this range are those that are accessed by different users at the same time, some for inserting into the file, some for updating records in the file and some for reading the records in the file.

Note: For more detail on how to perform the functions above and what the reports contain as well as how to interpret the data, please consult the PERFORMANCE TOOLS/400 GUIDE publication, SC41-4340-00.

DB2/400 SQL and Query Information Improvements

In V3R6, there have been several enhancements made to improve performance for SQL queries. In addition, more information is now available to help users analyze performance for all queries.

- 1. In V3R6, there is now more information provided to query users to help them analyze and improve the performance of their queries:
 - Index advisor messages have now been added to the messages that are generated in the joblog when running a query in a job in DEBUG mode. These messages will indicate how an index could be constructed that would be optimal for the performance of that query. Note that the information provided is generally most useful for queries that involve a single file or for the

primary file in a join query.

- There is a new database performance monitor function available via the STRDBMON (Start Database Monitor) command. This monitor will provide detailed information on all DB2/400 queries, such as CPU, I/O, elapsed time, description of the query, etc. The information is placed in a database file where it can be readily queried. The data provided by this function can provide valuable information for performance analysis of any DB2/400 query. For more information on the database monitor, refer to the *DB2 for OS/400 Database Programming* guide.
- 2. Support has been added to allow more types of join operations such as outer joins and exception joins. Although users could previously construct queries involving UNIONs to do outer joins, this was often cumbersome to do. The new SQL join syntax now gives users the ability to easily code these types of queries, often with significantly improved performance versus the previous alternative methods. For more information on the new syntax, refer to the *DB2 for OS/400 SQL Programming* guide, (SC41-4611).
- 3. The new join syntax for SQL now also gives users the ability to specify the order in which they want files to be joined. This can help improve performance for join queries where the optimizer is not choosing the optimal order in which to do the join. More information on this is available in the above mentioned SQL Programming guide.
- 4. Prior to V3R6, all SQL cursors within a given program/module operated under the commitment control level specified for that program when it was compiled. Now, however, a new WITH clause has been added to the SQL SELECT statement to allow specific cursors to run under the desired level of commitment control. For example, if an SQL program is running with a commit level of *ALL but there are read-only cursors in the program that do not need any commitment control, you can add the WITH NC clause to the SELECT statements to have these cursors run under a level of *NONE. Other levels that can be specified are UR (for *CHG), CS (for *CS) and RS (for *ALL). In many cases, a significant improvement in performance can be realized when this type of change is made to cursors that previously had been running under a more stringent level of commitment control than was necessary.
- 5. Support has been added to the ALTER TABLE SQL command that gives users the ability to easily add new fields and delete/change existing fields in any database file. Although users could do this in the past by deleting the old file and recreating it with a new format, this also meant that any views and indexes over the file had to be rebuilt as well, which could take a long time to complete. With the new support, the existing database file is copied to another file with the new format, and the indexes over the file are not rebuilt as long as no key fields are being altered. Although it still may take a while to do this copy, altering a database file's format with ALTER TABLE should be considerably faster than what the user had to do previously.
- 6. Prior to V3R6, a UNION ALL operation that did not specify an ORDER BY always generated a temporary file containing the results from each of the SELECT statements. This also meant that the ODP for this UNION was not reusable, which resulted in a full open and close each time the UNION was run. In V3R6, this type of operation now operates with 'live' data, i.e., the results from the second SELECT are not read until all the rows from the first SELECT have been read. This change will in most cases result in significantly improved response times for the first rows returned from the UNION since these rows can be returned immediately without having to wait for

the entire temporary file to be filled with results from all the SELECTs. Also, SQL is now able to make the ODP for this operation reusable, which also improves response time significantly.

- 7. In previous releases, the ODP for a cursor that contained a LIKE clause with a host variable mask was not reusable, which meant a full open was required each time the query was run. In V3R6, the ODP will be reusable if the value in the host variable mask is of the form 'XXXX%' and the NUMBER of constant characters in the mask stays exactly the same between each run of the query. In the case shown here, the contents of the XXXX constant part may be changed, but the number of constant characters (4) must remain the same and there cannot be anything else in the mask. If these rules are adhered to, users can significantly improve the performance of this type of SQL query.
- 8. Queries such as UNIONs, subqueries and joins may in some cases specify the same view in each SELECT specified in the query. Prior to V3R6, running these types of queries resulted in the view being evaluated multiple times, once for each time it was specified. In V3R6, the view in cases like this is evaluated only once and the results are used by each SELECT in the query. This change can result in a noticeable improvement in performance for this type of query, particularly if each evaluation of the view is costly.

DB2/400 SQL

DB2/400 Structured Query Language (SQL) support provides the user with an additional means of accessing data within an OS/400 relational database. This support provides several advantages in terms of flexibility, productivity and portability between various database platforms. However, prior to using SQL, users should also consider what level of performance they can expect when using this product. This section will provide general information on the performance of SQL to help users better determine what this level of performance will be.

This section is not intended to be a complete guide to SQL performance. For many users, other items of interest will include the SQL optimizer, performance tips and techniques, database design, etc. It is important that users take advantage of SQL performance tips and techniques as much as possible when writing an application using SQL. In particular, it is very important to properly construct and use indexes to provide the best overall performance for SQL. This, along with many other tips and techniques, can be found in the sources listed in "Additional Sources of Information" on page 48.

In V3R6, enhancements have been made to SQL that will help many users obtain an overall performance improvement. For a description of these enhancements, refer to "DB2/400 SQL and Query Information Improvements" on page 44.

Performance of SQL versus Native DB

When current users of AS/400 native language I/O (i.e., COBOL/400, RPG/400, etc.) are considering using SQL in their applications, one key item that needs to be considered is what level of difference in performance to expect when making this change.

Note that this section will not provide detailed information on the difference in performance between native and SQL for specific types of I/O operations. However, this type of information can be found in Chapter 5 of the document entitled "SQL/400 - A Guide for Implementation" (GG24-3321-01). Also, Chapter 3 of

the second version of this document (GG24-3321-02) contains a section entitled "When to Use SQL" that provides additional information and guidelines on this subject.

It is difficult to predict how SQL will compare to native DB access for a given application. Generally, SQL will use anywhere from 10-30% more CPU than native, although this may vary considerably depending on the type of operations being done. For example, SQL shows considerably more overhead than native when operating on one record at a time, such as an OPEN-FETCH-UPDATE WHERE CURRENT OF-CLOSE sequence. However, for more complex operations or for operations involving a lower number of SQL statements, SQL will in many cases show relatively equal or better performance levels than native.

Note that the difference between SQL and native performance is usually in terms of CPU (the amount of I/O that occurs is generally about the same for native and SQL for similar functions). Note, however, that on systems where the CPU utilization has not reached the knee of the performance curve, a difference in CPU of 10-30% per transaction will not result in a large difference in response time. Beyond the knee, however, the response time difference may grow considerably.

Other Performance Notes

- Generally, the use of SQL will result in significantly increased memory requirements when compared to similar native DB operations. This is mainly due to the additional internal structures and program automatic storage required by SQL to maintain optimum performance levels, as well as the fact that SQL cannot share ODPs across or within applications as native DB can. However, it is important to remember that the extra memory required can vary widely from application to application and is mostly dependent on the complexity of the application. Simple applications involving only a small number of I/O operations may require little additional memory for SQL, but for complex applications involving many I/O operations the difference in memory requirements between native and SQL can be significant. When using SQL, users should monitor memory utilization using the AS/400 Performance Tools to better determine if additional memory will be required. In addition, support is available through the Quicksizer tool on HONE to help users size their system for SQL.
- Some SQL users may notice an increase in the amount of auxiliary storage used once their applications begin running. The main reason for this is again largely due to the number and complexity of the ODPs and other internal structures that are maintained by SQL for each individual user for optimum performance. It is important to remember that this additional storage requirement is only temporary, i.e., when the user's job ends, the storage will return to normal levels. However, for systems where auxiliary storage usage is already high, some evaluation of the number of active SQL users and their storage requirements may be needed prior to any large scale implementation of a given SQL application.
- If memory and/or auxiliary storage is a concern, there are ways to reduce consumption of these resources for SQL applications. Following are some methods that can be used:
 - In a given SQL program, combine any duplicate or like SQL statements into one statement in an internal procedure, and then call that procedure as needed. This will help reduce the number of ODPs for that program.
 - If an internal procedure containing SQL statements is duplicated in several different SQL programs, the number of ODPs can be reduced by placing this internal procedure into a separate SQL program and then calling that program as needed. However, the user needs to be careful when

doing this to ensure that the ODPs in this common program are reusable across different invocations of the program. Also, since external calls will cause some performance degradation, this also needs to be considered prior to implementing this type of change.

Some user applications "pre-open" all their SQL ODPs in order to avoid full opens when the SQL statements are issued. Although this will provide good performance in many cases, there may be ODPs that are pre-opened but rarely used. If this is true, the user may want to be more selective about which ODPs are pre-opened and which are left to be opened when the SQL statement is first issued.

Additional Sources of Information

There are several other sources of information available that the user should obtain to gain a better understanding of SQL and OS/400 database operations, as well as understanding how to properly code SQL functions in order to optimize performance. Following is a list of these sources.

• SQL/400 - A Guide for Implementation, GG24-3321

There are three versions of this document currently available (01, 02 and 03). All versions contain key information and concepts users attempting to gain a better understanding of SQL operations.

This information is also available on HONE. Please refer to HONE item number RTA000011842.

• *Quicksizer* support for SQL

Available on HONE to help users size their system for using SQL.

• System Selection Guide

Available on HONE under the title "System Selection Guide". Hardcopy editions of the U.S. version and the worldwide version are also available.

- *DB2 for OS/400 SQL Programming*, SC41-4611
- DB2 for OS/400 Database Programming
- SQL/400 Programmer's Guide, SC41-9609
- SQL/400 Reference, SC41-9608
- AS/400 Database Guide, SC41-9659
- New Products Planning Information (NPPI), GA41-0007
- FTN broadcasts entitled "SQL Performance Technical Update"

There are several versions of this broadcast currently available on videotape. Each version covers SQL improvements made in each of the past several releases.

- Classes and workshops available to help users understand SQL programming and OS/400 database design and coding. One such recommended workshop is the "SQL Performance Workshop".
- SQL presentations from U.S. and European COMMON conferences.

Query Management

OS/400 Query Management (QM) is the AS/400 implementation of the SAA Common Programming Interface (CPI) Query. It provides a common method for accessing data and reporting the results from a relational database across the different platforms allowed by SAA. QM is also a very powerful and flexible reporting tool that provides users with the ability to design and format printed reports that result from the processing of a query. Queries can be included in programs written in RPG, COBOL, or C language and also can be run from within CL programs, giving programmers flexibility in how they set up the environment.

As a general rule, QM queries perform noticeably slower than similar functions that are done via AS/400 Query or embedded SQL because of the additional CPU used by QM (disk I/O characteristics are similar to those of AS/400 Query and SQL). For example, generating large reports via QM in most cases is significantly slower than AS/400 Query, and transaction processing with QM compares poorly with other alternatives such as static SQL. The exception to this rule is in processing summary-only functions such as AVG, COUNT, MIN, MAX and SUM. For this type of function, QM offers equal or better performance than AS/400 Query and performance levels similar to those when using SQL. Also, QM is optimized toward producing and displaying the first screen of output, so response times for this type of activity may be better than that for generating reports from QM.

QM should generally not be considered as an end-user tool. However, the SQL/400 Query Manager is available as an end-user type of interface for QM. Since it sits on top of QM and uses QM support, performance for this product will be similar to that of QM.

In general, customers who are considering the use of OS/400 Query Management need to weigh the functional advantages and flexibility (particularly in the area of portability across various SAA platforms) against the overall performance level of this product. Doing this will help decide if QM is a viable alternative to other AS/400 query products.

Referential Integrity

In a database user environment, there are frequent cases where the data in one file is dependent upon the data in another file. Without support from the database management system, each application program that updates, deletes or adds new records to the files must contain code that enforces the data dependency rules between the files. Referential Integrity (RI) is the mechanism supported by DB2/400 that offers its users the ability to enforce these rules without specifically coding them in their application(s). The data dependency rules are implemented as referential constraints via either CL commands or SQL statements that are available for adding, removing and changing these constraints.

For those customers that have implemented application checking to maintain integrity of data among files, there may be a noticeable performance gain when they change the application to use the referential integrity support. The amount of improvement depends on the extent of checking in the existing application. Also, the performance gain when using RI may be greater if the application currently uses SQL statements instead of HLL native database support to enforce data dependency rules.

When implementing RI constraints, customers need to consider which data dependencies are the most commonly enforced in their applications. The customer may then want to consider changing one or more of these dependencies to determine the level of performance improvement prior to a full scale implementation of all data dependencies via RI constraints.

Triggers

Trigger support for DB2/400 allows a user to define triggers (user written programs) to be called when records in a file are changed. Triggers can be used to enforce consistent implementation of business rules for database files without having to add the rule checking in all applications that are accessing the files. By doing this, when the business rules change, the user only has to change the trigger program.

There are three different types of events in the context of trigger programs: insert, update and delete. Separate triggers can be defined for each type of event. Triggers can also be defined to be called before or after the event occurs.

Generally, the impact to performance from applying triggers on the same system for files opened without commitment control is relatively low. However, when the file(s) are under commitment control, applying triggers can result in a significant impact to performance.

Triggers are particularly useful in a client server environment. By defining triggers on selected files on the server, the client application can cause synchronized, systematic update actions to related files on the server with a single request. Doing this can significantly reduce communications traffic and thus provide noticeably better performance both in terms of response time and CPU. This is true whether or not the file is under commitment control.

The following are performance tips to consider when using triggers support:

- Triggers are activated by an external call. The user needs to weigh the benefit of the trigger against the cost of the external call.
- If a trigger is going to be used, leave as much validation to the trigger program as possible.
- Avoid opening files in a trigger program under commitment control if the trigger program does not cause changes to commitable resources.
- Since trigger programs are called repeatedly, minimize the cost of program initialization and unneeded repeated actions. For example, the trigger program should not have to open and close a file every time it is called. If possible, design the trigger program so that the files are opened during the first call and stay open throughout. To accomplish this, avoid SETON LR in RPG, STOP RUN in COBOL and exit() in C.
- If the trigger program opens a file multiple times (perhaps in a program which it calls), make use of shared opens whenever possible.
- If the trigger program is written for the Integrated Language Environment (ILE), make sure it uses the caller's activation group. Having to start a new activation group every time the time the trigger

program is called is very costly.

• If the trigger program uses SQL statements, it should be optimized such that SQL makes use of reusable ODPs.

In conclusion, the use of triggers can help enforce business rules for user applications and can possibly help improve overall system performance, particularly in the case of applying changes to remote systems. However, some care needs to be used in designing triggers for good performance, particularly in the cases where commitment control is involved.

System-Managed Access-Path Protection (SMAPP)

Description

System-Managed Access-Path Protection (SMAPP) offers system monitoring of potential access path rebuild time and automatically starts and stops journaling of system selected access paths dynamically in order to meet a specified access path recovery time.

The default system wide access path recovery time for SMAPP is 150 minutes. This means that SMAPP protects the system so that there will be no more than 150 minutes of access path rebuild time during an IPL after an abnormal termination. Users can easily alter this value through the EDTRCYAP (Edit Recovery for Access Paths) or CHGRCYAP (Change Recovery for Access Paths) commands. SMAPP takes over the responsibility of providing the necessary amount of protection. No user intervention is required as SMAPP will manage the entire journal environment.

For systems with user auxiliary storage pools (ASPs), the recovery time can be specified for each ASP rather than one number for the entire system. This granularity allows the users to specify recovery time according to the criticality of the data on these ASPs. However, it is not recommended to specify target access path recovery times for both the entire system and individual ASPs.

For more information on SMAPP, see the Backup and Recovery - Advanced Book (SC41-3305).

SMAPP Impacts on Overall System Performance

The overhead of SMAPP varies from system to system and application to application due to the number of variables involved. For most customers, the default value of 150 minutes will minimize the performance impact while at the same time providing a reasonable and predictable recovery time and protection for key access paths. For many environments, even 60 minutes of IPL recovery time will have negligible overhead. Although SMAPP may start journaling access paths, the underlying SMAPP support is designed to be much cheaper in terms of performance than explicit journaling support.

Note that as the target access path recovery time is lowered, the performance impact from SMAPP will increase. You should balance your recovery time requirements against the system resources required by SMAPP.

Although the default level of SMAPP protection will be sufficient for most customers, some customers will need a different level of protection. The important variables are the number of key changes and the number of unprotected access paths. For those users who have experienced abnormal IPL access path recovery longer than 150 minutes it is advisable to experiment by varying the amount of protection. Too much

protection causes undue CPU consumption whereas too little protection causes undesirable IPL delay. Customers may need to decide on an optimum SMAPP setting by understanding their system requirements and experimenting to find what value meets these requirements.

There is some help for those who want to experiment. The component report produced by the licensed program *Performance Tools/400* has a database journaling summary. It has information that can help explain the effects of various SMAPP settings. This information is also available to all customers without this licensed program except it takes a little work to query the information (see the chapter titled *Collecting Performance Data* in the *Work Management Guide*).

Users may also experience more DASD usage if they are explicitly journaling their physical files and SMAPP starts journaling for the access paths to the same user journal. However, this increase may be lessened by using the RCVSIZOPT(*RMVINTENT) option on the CRTJRN or CHGJRN command. This will cause the system to remove internal entries used only for IPL recovery when they are no longer needed.

There will be some customer environments (such as those having a tight batch window) where no additional performance overhead can be tolerated. For these environments, it recommended that the SMAPP setting be changed to a much higher number or *NONE prior to the batch window and then changed back to the default/chosen value during transaction-heavy hours.

If ANY overhead at all cannot be tolerated, SMAPP can be turned off completely (special value *OFF). In this mode, there is no performance overhead, but there is also no idea of how exposed the system is. Also, to turn SMAPP back on, the system must be in a restricted state. Therefore, it is not advisable to turn SMAPP *OFF. The differences between SMAPP *NONE and SMAPP *OFF are:

- SMAPP *NONE allows SMAPP to monitor the system exposure without journaling access paths.
- You do not have to be in a restricted state to change from SMAPP *NONE to any other setting.

Miscellaneous Notes

- 1. SMAPP has no performance impact when you run applications with no access paths or those that do not make any key changes.
- 2. If SMAPP has a noticeable impact to performance, it will generally be in terms of increased CPU utilization and/or increased asynchronous IO activity. In most cases, SMAPP will have little effect on the the amount of synchronous IO.
- 3. The system starts to journal ALL access paths when SMAPP is set at *MIN (minimum access rebuild time during IPL or maximum protection). In some environments, the overhead of *MIN can result in a significant impact to overall system performance. For this reason, *MIN is not a recommended setting. If you have several small access paths that have many key changes, you are better off paying the small price of rebuilding them in the IPL following an abnormal termination (which is not frequent) than paying the runtime overhead of maximum SMAPP protection.
- 4. SMAPP and explicit journaling (of physical files and/or access paths) can coexist and are compatible with each other.

5. If SMAPP decides to journal an access path for a physical file that is currently not being explicitly journaled, SMAPP must journal both the physical file and the access path. The impact from this change can be noticeable to an application's performance. However, if SMAPP also decides to journal more access paths for the physical file, the added cost of journaling each additional access path will be less than the impact from journaling the first access path.

Journaling and Commitment Control

This section provides performance information and recommendations for DB2/400 journaling and commitment control.

Journaling

The primary purpose of journal management is to provide a method to recover database files. Additional uses related to performance include the use of journaling to decrease the time required to back up database files and the use of access path journaling for a potentially large reduction in the length of abnormal IPLs. For more information on the uses and management of journals, refer to the *AS/400 Backup and Recovery Guide*.

• The addition of journaling to an application will impact performance in terms of both CPU and I/O as the application changes to the journaled file(s) are entered into the journal. Also, the job that is making the changes to the file must wait for the journal I/O to be written to disk, so response time will in many cases be affected as well.

Journaling impacts the performance of each job differently, depending largely on the amount of database writes being done. Applications doing a large number of writes to a journaled file will most likely show a significant degradation both in CPU and response time while an application doing only a limited number of writes to the file may show only a small impact.

• The impact to performance from adding journaling can be reduced by locating the journal receiver on a separate user ASP. Doing this will generally reduce the seek time required to access the disk arms for journal I/O which will in turn help reduce the impact to end user response time. It will also lessen the impact to the disk arms located on the system ASP.

When using a separate user ASP for journal receivers, it is important to consider the number of disk actuators available in the ASP. Customer environments with heavily used journal receivers located in a user ASP that consists of a single disk actuator may actually reach a limit to performance because of the high usage of this single actuator. In this case, it would be better to have multiple disk actuators available in the user ASP so that DB2/400 journaling support can interleave journal entries over the multiple actuators, thus reducing contention for any one single disk arm. Doing this may result in an improvement in response time and in overall system throughput. However, it is important to note that although adding an actuator may provide a significant improvement in performance, each additional actuator added beyond this will improve performance to a lesser degree. Once the utilization of the actuators is low, adding more actuators will not improve performance.

Having two or more journal receivers located on the same user ASP and having them in use at the same time may not take full advantage of the performance gains seen by isolating a single journal receiver on the User ASP since the seek distance on the actuator increases as the journal entries are written to the

two receivers.

Tracked asynchronous I/O is used to write the journal entries to disk. The use of this type of I/O allows
the journal support to determine on a process by process basis, which processes need to wait for the
I/O to complete and which are allowed to continue. However, by using tracked asynchronous I/O, all
I/O operations to a journal receiver now appear in performance reports as asynchronous even though
the process may actually be waiting for the I/O operation to complete. This could cause the Capacity
Planning tools to recommend a smaller configuration than is necessary. This should beconsidered if a
measured profile is created for purposes of future system capacity planning.

Commitment Control

Commitment control is an extension to the journal function that allows users to ensure that all changes to a transaction are either all complete or, if not complete, can be easily backed out. The use of commitment control adds two more journal entries, one at the beginning of the committed transaction and one at the end, resulting in additional CPU and I/O overhead. In addition, the time that record level locks are held increases with the use of commitment control. Because of this additional overhead and possible additional record lock contention, adding commitment control will in many cases result in a noticeable degradation in performance for an application that is currently doing journaling.

- There are instances where adding commitment control can result in improved response times for an application doing journaling. As stated before, journaling alone means that the journal entries for changes to the file are written synchronously to disk. However, under commitment control, most journal entries are written to disk asynchronously. Only the final journal entry of the commit cycle (along with any entries of the cycle that have not yet been written to disk) are written synchronously. Because of this, applications may no longer have to wait for each journaled change to be written, which can result in reduced response times. The amount of improvement will depend mainly on the number of journal entries within the commit cycles the more entries per cycle, the greater the potential for improving response time over journaling alone. For example, adding commitment control to a dedicated batch job that is currently doing journaling could potentially improve the job run time if there are a large number of changes to the physical files being journaled.
- It is important to remember that the potential for improving response time by adding commitment control is also largely affected by overall system resource utilization. Environments that are showing high CPU or disk utilization or have constrained memory will in most cases show a degradation in performance from adding commitment control because of the additional CPU and I/O required. Also, adding commitment control can result in record level lock contention between jobs, which can also affect response time. Given the number of variables involved, a test run is highly recommended prior to adding commitment control for the purpose of improving performance in a production environment.

Date/Time Fields

The support for date and time fields in DB2/400 provides a number of advantages for the end user:

• Programmer productivity may be improved when an application requires calculations on date or time fields. New functions can be added more easily.

- Since the date and time data is stored in an internal format and converted on retrieval, the same underlying data can be viewed in different formats based on the needs of the application.
- Because the internal format reflects the sequential nature of time, it can be easily used to sort data in terms of sequence. For example, if a file currently contains a date in MMDDYY format, special application processing is required to sort it in YYMMDD sequence. This application processing is not needed when the date is stored in internal format.
- Some applications may achieve small savings in file size and DASD requirements since the internal formats are generally smaller than external formats.

The use of DB2/400 date/time support in many cases will result in additional CPU resource being used. Generally, the increase will be less than 10% but is dependent upon the number of calculations and the number of date and time fields being accessed. Time fields will usually show minimal impact while date and timestamp data types may show more of an effect on performance.

Note that in terms of performance, DB2/400 date/time support is generally better than or equal to other generalized routines that support many different date/time formats. However, when compared to date/time routines that handle only very specific date/time formats, DB2/400 date/time support may have higher CPU requirements.

When using date/time support in products such as AS/400 Query and Query Management, the amount of additional CPU required will vary. In many instances, the impact will be minimal and may even show a small reduction in CPU versus previous methods of providing this type of support. For example, report breaks on date fields under AS/400 Query will in many cases provide comparable performance to using packed data for dates. However, there are certain cases where the use of date/time support can result in significant performance overhead:

- When replacing the use of zoned decimal data for dates
- When adding a result field calculations to a query (such as adding 90 days to a date)
- Report breaks on date fields under Query Management (compared to the use of packed data for dates)

Overall, DB2/400 date/time support can provide many functional advantages to user applications without a significant impact to performance. However, the user should exercise some caution when implementing this support in order to minimize this impact.

Null Values

DB2/400 provides support for the use of null values in any field in any file. For a more detailed description of null value support, refer to the SQL/400 Reference or the SQL/400 Programmers Guide.

The performance impact from using null value support will vary depending on the number of fields declared as null capable and on the number of records being accessed. For example, when a user even changes only one field in a file to be null capable, there will be a slight increase in the CPU resource required to either insert records into or read records from this file. The amount of the increase should be about the same whether or not the null capable field actually contains null values. Also, as the number of null capable fields in a given file record format increases, the CPU required to process each record will also

increase. For operations such as AS/400 Query, Query Management and SQL/400 queries that select all the fields from a large number of records, the impact of adding null capable fields to the file can be significant in terms of increased CPU.

Because of the potential impact, users need to be somewhat careful in what files null capable fields will be used and in deciding how many fields will be null capable. Although null capable fields do provide good functional advantages, performance also needs to be considered prior to using this support.

CCSID Support

CCSID (Coded Character Set Identification) enhancements support the dynamic conversion of data from one language to another. The support allows jobs, files, and fields within files to be tagged with an identification of the code page currently being used. For a more detailed description of this support refer to the AS/400 National Language Support Planning Guide.

The main effect to performance from CCSID support is from the character data conversion required when either the CCSID of the job and the file/field do not match or when either of these CCSID values is not set to 65535. The amount of additional CPU required for this conversion will vary somewhat depending on the amount of character data that needs to be converted. Since the impact of this conversion can be significant to normal database operations, users should exercise some caution when implementing this function. For example, it may be best to consider doing CCSID conversion only on fields that need the conversion done instead of all character data in the given database file.

Sort Sequence

DB2/400 sort sequence support provides application developers and end users with an easy method of producing sorted data for a particular language or culture. A set of unique and shared sort sequence tables are included on the AS/400. Developers can refer to sort sequences when creating applications using database, Query/400, RPG, COBOL, C, and ILE/C compilers, as well as SQL precompilers.

The performance of sort sequence support should be compared to the alternative methods that users have available on the AS/400. For example, users who desire to use a different sorting sequence in QUERY/400 queries can create a translation table and then specify this translation table in the "select alternate collating sequence" option in QUERY/400. However, comparisons of these two methods show that sort sequence support will provide a noticeable improvement in performance (ranging from 5-40%) versus using the translation table method.

Users who would like to learn more about sort sequences should refer to the *National Language Support Planning Guide*.

Variable Length Fields

Variable length field support allows a user to define any number of fields in a file as variable length, thus potentially reducing the number of bytes that need to be stored for a particular field.

Description

Variable length field support on the AS/400 has been implemented with a spill area, thus creating two possible situations: the non-spill case and the spill case. With this implementation, when the data

overflows, all of the data is stored in the spill portion. An example would be a variable length field that is defined as having a maximum length of 50 bytes and an allocated length of 20 bytes. In other words, it is expected that the majority of entries in this field will be 20 bytes or less and occasionally there will be a longer entry up to 50 bytes in length. When inserting an entry that has a length of 20 bytes or less that entry will be inserted into the allocated part of the field. This is an example of a non-spill case. However, if an entry is inserted that is, for example, 35 bytes long, all 35 bytes will go into the spill area.

To create the variable length field just described, use the following SQL/400 statement:

```
CREATE TABLE library/table-name
(field VARCHAR(50) ALLOCATE(20) NOT NULL)
```

In this particular example the field was created with the NOT NULL option. The other two options are NULL and NOT NULL WITH DEFAULT. Refer to the NULLS section in the SQL/400 Reference Guide to determine which NULLS option would be best for your use. Also, for additional information on variable length field support, refer to either the SQL/400 Reference Guide or the SQL/400 Programmer's Guide.

Performance Expectations

- Variable length field support, when used correctly, can provide performance improvements in many environments. The savings in I/O when processing a variable length field can be significant. The biggest performance gains that will be obtained from using variable length fields are for description or comment types of fields that are converted to variable length. However, because there is additional overhead associated with accessing the spill area, it is generally not a good idea to convert a field to variable length if the majority (70-100%) of the records would have data in this area. To avoid this problem, design the variable length field(s) with the proper allocation length so that the amount of data in the spill area stays below the 60% range. This will also prevent a potential waste of space with the variable length implementation.
- Another potential savings from the use of variable length fields is in DASD space. This is particularly true in implementations where there is a large difference between the ALLOCATE and the VARCHAR attributes AND the amount of spill data is below 60%. Also, by minimizing the size of the file, the performance of operations such as CPYF (Copy File) will also be improved.
- When using a variable length field as a join field, the impact to performance for the join will depend on the number of records returned and the amount of data that spills. For a join field that contains a low percentage of spill data and which already has an index built over it that can be used in the join, a user would most likely find the performance acceptable. However, if an index must be built and/or the field contains a large amount of overflow, a performance problem will likely occur when the join is processed.
- Because of the extra processing that is required for variable length fields, it is not a good idea to convert every field in a file to variable length. This is particularly true for fields that are part of an index key. Accessing records via a variable length key field is noticeably slower than via a fixed length key field. Also, index builds over variable length fields will be noticeably slower than over fixed length fields.
- When accessing a file that contains variable length fields through a high-level language such as COBOL, the variable that the field is read into must be defined as variable or of a varying length. If

this is not done, the data that is read in to the fixed length variable will be treated as fixed length. If the variable is defined as PIC X(40) and only 25 bytes of data is read in, the remaining 15 bytes will be space filled. The value in that variable will now contain 40 bytes. The following COBOL example shows how to declare the receiving variable as a variable length variable:

```
01 DESCR.

49 DESCR-LEN PIC S9(4) COMP-4.

49 DESCRIPTION PIC X(40).

EXEC SQL

FETCH C1 INTO DESCR

END-EXEC.
```

For more detail about the vary-length character string, refer to the SQL/400 Programmer's Guide.

The above point is also true when using a high-level language to insert values into a variable length field. The variable that contains the value to be inserted must be declared as variable or varying. A PL/I example follows:

```
DCL FLD1 CHAR(40) VARYING;
FLD1 = XYZ Company;
EXEC SQL
INSERT INTO library/file VALUES
("001453", FLD1, ...);
```

Having defined FLD1 as VARYING will, for this example, insert a data string of 11 bytes into the field corresponding with FLD1 in this file. If variable FLD1 had not been defined as VARYING, a data string of 40 bytes would be inserted into the corresponding field. For additional information on the VARYING attribute, refer to the PL/I User's Guide and Reference.

• In summary, the proper implementation and use of DB2/400 variable length field support can help provide overall improvements in both function and performance for certain types of database files. However, the amount of improvement can be greatly impacted if the new support is not used correctly, so users need to take care when implementing this function.

Reuse Deleted Record Space

Description of Function

This section discusses the support for reuse of deleted record space. This database support provides the customer a way of placing newly-added records into previously deleted record spaces in physical files. This function should reduce the requirement for periodic physical file reorganizations to reclaim deleted record space. File reorganization can be a very time consuming process depending on the size of the file and the number of indexes over it. To activate the reuse function, set the Reuse deleted records (REUSEDLT) parameter to *YES on the CRTPF (Create Physical File) or CHGPF (Change Physical File) commands. The default value when creating a file is *NO (do not re-use).

Comparison to Normal Inserts

Inserts into deleted record spaces are handled differently than normal inserts and have different performance characteristics. For normal inserts into a physical file, the database support will find the end of the file and seize it once for exclusive use for the subsequent adds. Added records will be written in blocks at the end of the file. The size of the blocks written will be determined by the default block size or by the size specified using an Over-ride Database File (OVRDBF) command. The SEQ(*YES number of records) parameter can be used to set the block size.

In contrast, when re-use is active, the database support will process the added record more like an update operation than an add operation. The database support will maintain a bit map to keep track of deleted records and to provide fast access to them. Before a record can be added, the database support must use the bit-map to find the next available deleted record space, read the page containing the deleted record entry into storage, and seize the deleted record to allow replacement with the added record. Lastly, the added records are blocked as much as permissible and then written to the file.

To summarize, additional CPU processing will be required when re-use is active to find the deleted records, perform record level seizes and maintain the bit-map of deleted records. Also, there may be some additional disk IO required to read in the deleted records prior to updating them. However, this extra overhead is generally less than the overhead associated with a sequential update operation.

Performance Expectations

The impact to performance from implementing the reuse deleted records function will vary depending on the type of operation being done. Following is a summary of how this function will affect performance for various scenarios:

- When blocking was not specified, re-use was slightly faster or equivalent to the normal insert application. This is due to the fact that reuse by default blocks up records for disk IOs as much as possible.
- Increasing the number of indexes over a file will cause degradation for all insert operations, regardless of whether reuse is used or not. However, with reuse activated, the degradation to insert operations from each additional index is generally higher than for normal inserts.
- The RGZPFM (Reorganize Physical File Member) command can run for a long period of time, depending on the number of records in the file and the number of indexes over the file. Even though activating the reuse function may cause some performance degradation, it may be justified when considering reorganization costs to reclaim deleted record space.
- The reuse function can always be de-activated if the customer encounters a critical time window where no degradation is permissible. The cost of activating/de-activating reuse is relatively low in most cases.
- Because the reuse function can lead to smaller sized files, the performance of some applications may actually improve, especially in cases where sequential non-keyed processing of a large portion of the file(s) is taking place.

DB2/SMP Feature

Introduction

The symmetrical multiprocessing (SMP) feature provides additional query optimization algorithms for retrieving data. In addition, the DB2/SMP feature provides application transparent support for parallel query operations on a single tightly-coupled multi-processor AS/400 system (shared memory and disk). The database manager can automatically activate parallel query processing in order to engage one or more system processors to work simultaneously on a single query. The response time can be dramatically improved when a processor bound query is executed in parallel on multiple processors. The purpose of this section is to:

- Introduce new query optimization algorithms available with the DB2/SMP feature.
- Briefly discuss decision support (DSS) queries which will realize the most benefit with the SMP feature.
- Provide guidance to help estimate DSS query capacity on various AS/400 systems.

New Query Optimization Algorithms

The DB2/SMP feature provides the following new query optimization algorithms:

• Parallel table scan

Provides parallel operations for queries requiring a sequential scan of the entire table. Multiple tasks are used to scan the same table concurrently. Each task will perform selection and column processing on a table partition and return selected records to the requester. The response time improvement for a parallel table scan scales closely to the number of processors participating. For example, the response time for a table scan can be up to 4 times faster when run in parallel on a 4-way processor.

• Index only access (parallel and non-parallel)

Provides performance improvement by extracting a query answer from an index rather than performing random I/Os against a physical table. For this to happen, all columns that are referenced in a query must exist within an index. Response time improvements can be up to 5 times faster for some queries.

• Parallel key selection

Provides parallel index operations for key selection. Multiple tasks are used to scan the same index concurrently. Each task will search a different key range and selected records are returned to the requester.

• Hashing algorithms

Provides an optimization alternative for group by and some join queries. This method avoids having to utilize an index and therefore avoids having to perform random I/Os to retrieve the results. Instead, a temporary partitioned hash table can be used. This table can be processed by large and efficient sequential I/Os and often utilizing parallel table scan to provide the results. Response time improvements for group by queries can be up 6 times better and some joins can be up to 25 times improved (4 to 10 times is more typical).

The SMP feature was available V3R1 on AS/400 IMPI models and became available V3R7 for AS/400 RISC models. For more information on the SMP feature and the new algorithms, see *TNL SN41-3680 to SC41-3611-00*.

Decision Support Queries

The SMP feature is most useful when running decision support (DSS) queries. DSS queries generally give answers to critical business questions tend to have the following characteristics:

- examine large volumes of data
- are far more complex than most OLTP transactions
- are highly CPU intensive
- includes multiple order joins, summarizations and groupings

DSS queries tend to be long running and can utilize much of the system resources such as processor capacity (CPU) and disk. For example, it is not unusual for DSS queries to have a response time longer than 20 seconds. In fact, complex DSS queries may run an hour or longer. The CPU required to run a DSS query can easily be 100 times greater than the CPU required for a typical OLTP transaction. Thus, it is very important to choose the right AS/400 system for your DSS query and data warehousing needs.

SMP Performance Summary

The SMP feature provides performance improvement for query response times. The overall response time for a set of DSS queries run serially at a single work station may improve 25 to 58 percent when SMP support is enabled. The amount of improvement will depend in part on the number of processors participating in each query execution and the optimization algorithms used to implement the query. Some individual queries can see significantly larger gains. Queries that are able to utilize the new hash join algorithm may see up to a 25 times improvement in query response time. In addition, query throughput may improve 18 to 25 percent because the new optimization algorithms require less CPU resource. The new hashing algorithms also dramatically reduce the number of disk IOs.

Capacity Planning

The Capacity Planning sections contain the following information:

- Initial system sizing recommendations for data warehouses
- Detailed capacity planning information for various AS/400 models. This information will be useful when you are able to determine a customer's average DSS query response time and want to compare running a query workload on other AS/400 models or with the SMP feature enabled.
- Capacity planning tips

System Sizing Recommendations for Data Warehouses

The following table gives some high-level guidance for choosing the AS/400 system for Data Warehouses based on the size of the database and/or the maximum number of concurrent users.

Table 16. System Sizing Recommendations for AS/400 Data Warehouses						
System	Maximum Data in Gigabytes	Maximum Number of Concurrent Users				
40S	15	15				
50S 2120	220	20				
50S 2121	220	30				
538 2154	350	40				
538 2155	350	65				
538 2156	350	350 125				

Note:

1. The maximum amount of main storage exists on each system.

2. Query workloads are assumed to be comprised of the following query mixture:

• Simple queries (no joins or group by aggregation) 80%

- Medium queries (2-way joins, group by aggregation) 15%
- Complex queries (union, subselects) 5%

3. Simple query workloads may also include the use of any Multi-Dimensional Database product.

4. If your database size is greater than 350 Gigabytes or the number of concurrent users is greater than 125, you will require a multisystem implementation which may include the DB2 Multisystem feature. DB2 Multisystem provides the capability of horizontally partitioning a table across multiple systems and running a single query in parallel. Performance is improved due to parallel operations and because of the table partitioning. Each system needs only scan a fraction of the entire database when a query is run.

Capacity Planning based upon Average Query Response Time

This section provides some guidance to help you estimate DSS query capacity on various AS/400 systems, with and without the SMP feature enabled. The chart was developed based upon studying the results of various customer and synthetic DSS workloads. The workloads contained various sized files ranging from 25 records up to 100 million records. A broad range of DSS support queries, from simple to complex were measured. Queries that utilized joins, group by, and summarizations were commonplace. The database structure, the index structure, and the query syntax are all assumed to be optimal. The SMP numbers in the the chart show a range of performance based upon an estimate of:

- the percentage of the DSS queries that might be helped by SMP
- the query benefit provided by parallelism

How to use the chart

Calculating the capacity for DSS query workloads can be difficult due to vast variability of the queries. The capacity chart uses an average query response time that might be observed in a customer environment over a long period of time such as a day. Obviously, during this time period there will be great variance in terms of complexity of the queries, the size of the tables queried, the response times of the individual queries, and the load put on the AS/400 system.

For the capacity chart, we have used 180 seconds to represent a customer's average query time. An average of 180 seconds would indicate a majority of simple DSS queries being executed during the time window. If this does not accurately reflect your customer's environment, you can estimate new average response times and system capacities by performing the calculations that follow the capacity chart:

Table 17. System Capacity Planning - DSS Queries							
Model	CPUs	Avg Response Time in Seconds		Capacity in Queries/Hour			
		No SMP	SMP Range	No SMP	SMP Range		
30S 2411	1	180		40			
30S 2412	2	173	93-130	86	105-114		
E95	4	188	78-117	148	179-195		
F97 & 320	4	154	64-96	225	272-297		
50S 2121	1	100		98			
53S 2154	1	86		159			
53S 2155	2	86	46-64	244	296-324		
53S 2156	4	86	36-54	372	451-493		
Nata							

Note:

1. The average CPU reduction when SMP is enabled 18-25%.

2. Capacity numbers are based on 100% CPU utilization and assume that the system is dedicated to query processing.

3. Information in chart based on assumptions listed in the next section.

You can estimate new average response times and system capacities by performing the following calculations against the values in the capacity chart:

- Determine the customers average response time
- Compute the following query response time ratio: ratio = customer's average response time/average response time from table
- Multiply all response times by the ratio to get new response times
- Divide all capacities by the ratio to get new capacities
- For example, if the customer's average response time is 38.5 seconds on an F97 processor, to calculate the new F97 values, perform the following calculations:

Current F97 row: F97 & 320 4 154 64-96 225 272-297 Ratio = 38.5/154 = .25Response times Capacities _____ 154(*.25) = 38.5225/.25 = 900 No SMP 64(*.25) = 16.0272/.25 = 1088SMP 96(*.25) = 24.0297/.25 = 1188SMP F97 & 320 4 38.5 16-24 900 1088-1188 New F97 row: New 53S 2156 row: 53S 2156 4 21.5 9-13.5 1488 1804-1972

Capacity Planning Assumptions

The following assumptions were used to help generate the capacity chart:

- 1. DSS query workloads can be characterized by an average response time. The average response time will increase as the size of the customer's database size increases.
- 2. Given all of a customer's DSS queries, typically 50%-70% of the queries will utilize the SMP support.
- 3. For queries that utilize SMP, the response time will scale relative to the number of CPUs. The scaling range equals 1 to 1.5 times the number of CPUs involved in the query execution. For example, on a 4-way CPU system, response time will be 1/4 to 1/6 the time compared to executing the query on just one of the processors.
- 4. Hash group by/join algorithms will be utilized in about 70% of all queries that can utilize the SMP support. About 50-75% of the IOs will be eliminated when the new hashing algorithms are used.

- 5. Table scans will be utilized in about 10% of all the DSS queries.
- 6. For SMP queries, CPU consumption will decrease up to 35% due to the new optimization algorithms and because of the reduction in disk Ios.
- 7. A DSS query workload will utilize at least 50% of the system processor capacity when run on an AS/400 30S 2411.

Capacity Planning Tips

Here are some suggestions that may improve your DSS query performance when utilizing the SMP support:

- Add additional memory. 20%-25% of the active database should reside in main memory.
- Add additional disks and limit the the number of disks per controller to 8 if possible. This is especially true if you are using 9337 DASD and the 6501 DASD IOP, as this will cause more efficient use of active memory.
- Utilize fast DASD (6503,6506,and 6507) and DASD IOPs (6502 and 6512). Spread the IOPs evenly among the system busses.
- Ensure that the database is spread evenly over multiple DASD arms. Installing DASD that is all the same size helps ensure even spreading.
- For smaller database sizes (< 30GB), you should have 2-3 DASD arms per CPU to get good performance.
- Utilize RISC hardware. RISC systems have faster system busses with larger bandwidths than those found on IMPI systems. In addition, disk IO sizes are larger which will results in fewer disk Ios.
- Be sure that there is enough space on the the system auxiliary storage pool (ASP) to allow the database manager to create temporary files for query execution. Do not exceed 70% capacity on the system ASP.
- Under a heavy system load, limit the amount of query parallelism. The degree of parallel activity can be controlled by the user via the CHGSYSVAL (parm QQRYDEGREE) and CHGQRYA (parm Degree) CL commands.

DB2 Multisystem for OS/400

DB2 Multisystem for OS/400 offers customers the ability to distribute large databases across multiple AS/400s in order to gain nearly unlimited scalability and improved performance for many large query operations. The multiple AS/400s are coupled together in a shared-nothing cluster where each system uses its own main memory and disk storage. Once a database is properly partitioned among the multiple nodes in the cluster, access to the database files is seamless and transparent to the applications and users that reference the database. To the users, the partitioned files still behave as though they were local to their system.

This section will provide information on what level of performance improvements to expect from DB2 Multisystem as well as tips and techniques on how to install and use this product for optimal performance.

However, this section should not be viewed as a complete guide to performance for DB2 Multisystem. It is recommended that in addition to the information provided here, you should obtain the following documents to help understand more about both the key performance and functional aspects of this product.

• *DB2 Multisystem for OS/400*, SC41-3705-00

This document is an excellent overall reference for this product and contains several aspects of performance that will not be covered in this document, in particular some items on distributed query optimization and processing.

• Slash DB2/400 Query Time with Parallel Processing

This article (found in the April 1996 edition of the NEWS/400 magazine) helps explain key performance and functional concepts of DB2 Multisystem.

These documents and the information in this section assumes that you are familiar with nondistributed query performance on the AS/400 and that you have a good overall background in database concepts. Other documents that can help you with this information include:

- DB2 for OS/400 SQL Reference
- DB2 for OS/400 SQL Programming
- DB2 for OS/400 Database Programming
- CL Reference Guide

Planning for DB2 Multisystem

The most important aspect of obtaining optimal performance with DB2 Multisystem is to plan ahead for what data should be partitioned and how it should be partitioned. The main idea behind this planning is to ensure that the systems in the cluster run in parallel with each other as much as possible when processing distributed queries while keeping the amount of communications data traffic to a minimum. Following is a list of items to consider when planning for the use of distributed data via DB2 Multisystem.

- Avoid large amounts of data movement between systems. A distributed query often achieves optimal performance when it is able to divide the query among several nodes, with each node running its portion of the query on data that is local to that system and with a minimum number of accesses to remote data on other systems. Also, if a file that is heavily used for transaction processing is to be distributed, it should be done such that most of the database accesses are local since remote accesses may add significantly to response times.
- Choosing which files to partition is important. The largest improvements will be for queries on large files. Files that are primarily used for transaction processing and not much query processing are generally not good candidates for partitioning. Also, partitioning files with only a small number of records will generally not result in much improvement and may actually degrade performance due to the added communications overhead.
- Choose a partitioning key that has many different values. This will help ensure a more even distribution of the data across the multiple nodes. In addition, performance will be best if the partitioning key is a single field that is a simple data type.
- It is best to choose a partition key that consists of a field or fields whose values are not updated. Updates on partition keys are only allowed if the change to the field(s) in the key will not cause that record to be partitioned to a different node.
- If joins are often performed on multiple files using a single field, use that field as the partitioning key for those files. Also, the fields used for join processing should be of the same data type.
- It will be helpful to partition the database files based on how quickly each node can process its portion of the data when running distributed queries. For example, it may be better to place a larger amount of data on a large multiprocessor system than on a smaller single processor system. In addition, current normal utilization levels of other resources such as main memory, DASD and IOPs should be considered on each system in order to ensure that no one individual system becomes a bottleneck for distributed query performance. For information on how to customize your database partitioning, refer to the "DB2 Multisystem for OS/400" document mentioned above.
- For the best query performance involving distributed files, avoid the use of commitment control when possible. DB2 Multisystem uses two-phase commit, which can add a significant amount of overhead when running distributed queries.

In addition to these items, the document and article referenced above contain other key concepts that should be considered while planning your data distribution via DB2 Multisystem.

Performance During Data Distribution

Generally, partitioning large database files across multiple systems can be a long process during which the data in the files is unavailable. Following is a list of items that should be considered prior to actually partitioning the files to help reduce the time this process may take.

- The use of Opticonnect will result in significantly better distribution times than using other alternatives such as a 16Mbps Token Ring LAN. Opticonnect will also help improve performance for distributed queries that result in large amounts of data being moved from node to node to complete the query.
- There are basically two recommended methods of distributing data from a local system to a set of
 systems linked together with DB2 Multisystem. One method is to use the Change Physical File
 (CHGPF) command with the NODGRP and PTNKEY parameters. This command will need to be
 issued against each database file to be distributed. Any existing logical files for this file will also be
 rebuilt on a per node basis. The second method is to create a new physical file with the same data
 format as the original and with the node group and partition key specified (this can be done either via
 the Create Physical File (CRTPF) command or the SQL CREATE TABLE command), and then issue
 a Copy File (CPYF) command to copy the data from the original file to the new distributed file.
 Measurement results show that the performance of these methods is about equal.

Note that there is a faster and slower version of both the CHGPF and CPYF operations for distributing files. The faster version sends large buffers of records at a time while the slower version sends one record at a time. To see if the fast version is being done, look for occurrences of the CPC9203 message

in the joblog of the job doing the distribution, stating how many records were copied to each node. If these messages do not appear, the slower version is being used. The factors that influence which version is used are listed in more detail in the DB2 Multisystem for OS/400 document mentioned above.

- To help the distribution process, it may be best to keep the number of logical files to a minimum for the physical files that are being distributed. These logical files can then be built via the Create Logical File (CRTLF) or the SQL CREATE INDEX command at a later time, possibly in background batch jobs. This approach is generally faster than having the system maintain or build the indexes on each node as the physical data is distributed. However, you will have to issue the index builds separately and they will tend to cause high CPU utilization while they are occurring, so this must be considered as well. If you need certain key indexes to exist as soon as the data distribution is done, you should let the CPYF or CHGPF operations handle these for you.
- It will be to your benefit to avoid the use of commitment control or journaling while distributing database files. The use of these options will add significantly to the overall distribution time.
- The time for data distribution may also be helped by having several jobs running at the same time, each distributing a different file. Although this is best accomplished where the system doing the distribution is a multiprocessor system, this can also apply to single processor systems. The key to making this work is to avoid a bottleneck on a resource such as main memory, DASD, CPU or the communications lines or IOPs. It may be best to try this by adding one job at a time and monitoring system performance to see if any resources are becoming overutilized.

Distributed Query Performance

The performance of queries run over distributed data will in many cases improve significantly compared to the performance that had been achieved running these same queries on a single system. However, there also may be queries that show little or no performance gain, with some possibly showing degradation in performance. The following information should help determine what level of performance to expect when running queries over distributed data. Again, it is important to reference the above mentioned documents in conjunction with the information provided here in order to gain a more complete picture of distributed query performance.

- Use of the new ASYNCJ parameter on the Change Query Attributes (CHGQRYA) command is very important to achieving the best performance levels for distributed queries. The value specified for this parameter will greatly affect the response time for distributed queries by altering the degree of parallelism allowed as well as the amount of work done by the temporary result writer jobs. Note that this command needs to be issued on a per job basis as there is no global system level value that can be changed. For more information on the use of this command for distributed queries, refer to the *DB2 Multisystem for OS/400* document.
- There is now a distributed query optimizer that operates only on distributed queries. This optimizer determines what steps are necessary to efficiently run the distributed query and what nodes will process these individual steps. Local level optimization on each node is still handled by the previously existing query optimizer.
- The use of Opticonnect is recommended for the best overall performance of distributed queries. Although good planning will minimize the amount of communications overhead needed for many

distributed queries, there still will be a fair amount of cross-system data traffic in many DB2 Multisystem environments. Using Opticonnect will result in noticeably better response times for queries with a significant amount of cross-system data movement and will in general help reduce the communications overhead for users of DB2 Multisystem.

- Generally, the best performance gains from DB2 Multisystem will be for queries that exhibit the following characteristics:
 - The query processes a large number of records
 - The query can be divided such that subsets of the records it processes can be queried on multiple nodes in parallel
 - Each part of the divided query returns a small number of records to the coordinating system where the query originated

For queries that meet these criteria, performance can be expected to improve in nearly a linear progression with the number of systems involved in running the query. In addition, if any of the systems used are multiprocessor systems, the improvement on these nodes may also be multiplied by the enhancements provided by DB2 Symmetric Multiprocessing for OS/400 (SMP). For example, a query that had previously been run on a single processor and is now being run on three four-way systems could experience a run time that is one-twelfth of what it had been. Although this amount of improvement may not be realized in most queries, there will still be many queries that will experience large improvements in performance.

- Queries that read and process a small number of records may experience some level of performance improvement when running over distributed data, but the percentage of improvement will in many cases be much less than queries over large files. For queries of this type, the amount of improvement will often be a factor of the speed of the connection between the systems, and in some cases, this may cause the query to run longer than it had on a single system.
- Queries that read and process a large number of records but that also return a large number of records to the coordinating system will in many instances not experience the almost linear improvements mentioned above. In this case, the individual nodes may still be able to process a subset of the records efficiently, but the response time may be affected by how quickly the records in the individual answer sets can be transferred back to the coordinator and how quickly this system can receive and process them as well.
- The performance of join queries on distributed data is closely linked to how much data needs to be transferred between nodes to perform the join. The best performing join queries are where all of the corresponding records of the files being joined exist on the same node so that no data is moved to other nodes to perform the join. These types of joins should improve nearly linearly with the number of nodes, although this again depends on the amount of data that needs to be transferred back to the coordinating system and the additional processing that will be needed there. Other join operations that need to move data between nodes to do the join will vary widely in how much improvement is achieved, and in some cases, may end up with a significant degradation. For this reason, partitioning of commonly joined files needs to be planned such that the most common join operations end up moving only smaller amounts or no data between nodes. For a more detailed discussion on distributed join

performance, refer to Chapter 6 of the previously mentioned DB2 Multisystem for OS/400 document.

- Queries that specify selection criteria on a single file may end up doing all the processing of that query on a single node if the optimizer determines that all the records matching the criteria exist on that node. In this case, the amount of performance improvement for this type of query will vary depending on how quickly the system at that node can process the query and return the results to the coordinating system. However, there are certain restrictions that a query must meet in order to be directed to a single particular node. More information on this type of query can be found in Chapter 6 of the DB2 Multisystem document.
- For most distributed queries and in particular for queries involving ordering of data, it is best to specify the ALWCPYDTA(*OPTIMIZE) parameter on the Open Query File (OPNQRYF) and Start SQL (STRSQL) commands and also on the Create SQLxxx (CRTSQLxxx) commands. This option allows the optimizer the most flexibility in choosing what method to use (an index or a sort) to order the records on each node.
- To achieve the fastest retrieval of data from a distributed file, you can issue the Override Database File (OVRDBF) command with the DSTDTA(*BUFFERED) parameter specified. For more information on this option, refer to Chapter 5 of the previously mentioned DB2 Multisystem for OS/400 document.
 In addition to the above information, there are many other items to consider to understand distributed query optimization and how to obtain optimal performance levels when using DB2 Multisystem support. The following items (as well as many of the above) are covered in the above mentioned DB2 for Multisystem document.
- ORDER BY and GROUP BY operations
- Reusable and non-reusable ODPs
- Temporary result writers (new for DB2 Multisystem)
- Optimizer messages
- Changes to the Change Query Attributes (CHGQRYA) command

Chapter 5. Communications Performance

There are many factors that affect the performance of an AS/400 in a communications environment. This section discusses some of the common factors and offers guidance on how to achieve the best possible performance. Much of the information in this section was obtained as a result of analysis experience within the Rochester development laboratory. Many of the performance claims are based on supporting performance measurement and analysis with NetPerf and other performance workloads. In some cases, the actual performance data is included here to reinforce the performance claims and to demonstrate capacity characteristics.

V4R4 Communications Performance Highlights: The performance of the communications software infrastructure improved significantly in V4R4 based on performance improvements to Sockets, TCP/IP, and Ethernet. Many scenarios using this communications path reduced the CPU time required (for the communications-related portion) of a transaction by up to 50%. Also, the scalability of communications performance improved due to software enhancements to minimize contention. In addition, software enhancements for Ethernet support (with TCP/IP only) allow higher data rates and greater IOP capacities. New performance information has also been added for SSL and VPN. See the sections that follow for details.

Performance information, tips, and techniques for communications are listed in the following sections:

- 5.1 TCP/IP, Sockets, SSL, VPN, and FTP
- 5.2 APPC, APPN, ICF, CPI-C, and Anynet
- 5.3 LAN and WAN
- 5.4 Work Station Connectivity
- 5.5 Opti-connect for OS/400
- 5.6 NetPerf Workload Description

5.1 TCP/IP, Sockets, SSL, VPN, and FTP Performance Information

TCP/IP and Sockets Performance Information:

- In V4R3, TCP/IP and APPC generally provided similar levels of application performance.
- In V4R4, the performance of TCP/IP is superior to APPC. Significant network infrastructure software performance enhancements were implemented including: Optimization of Sockets APIs and re-implementing Sockets in SLIC, optimization of TCP/IP to improve performance and scalability, and a more efficient scheme for the Ethernet device driver to interface with the IOP/IOA. These changes improved performance by significantly reducing the CPU time required to process the communcations software.
- Measurements with the NetPerf workload demonstrates the significant performance improvement with V4R4 enhancements: The CPU time for the Request/Response scenario (client server like) was reduced by 40-50%. The CPU time for the Connect/Request/Response scenario (web server like) was reduced by 40-50%. The maximum transfer rate for the Streaming (large transfer) scenario using 100 Mbps Ethernet increased from 40 Mbps up to 90 Mbps. The NetPerf workload is defined later in this section.

- Always ensure that the entire communications network is configured optimally. The **maximum frame size parameter** (MAXFRAME on LIND) should be maximized. The **maximum transmission unit** (**MTU**) **size** parameter (CFGTCP command) for both the route and interface affect the actual size of the line flows and should be configured to *LIND. This means that there will be a one-to-one match between frames and MTUs.
- When transferring large amounts maximize the size of the application's send and receive size. This is the amount of data that the application transmits with a single sockets API. Being sockets does not block up multiple applications sends, it is important to limit the number of interactions. Block in the application if possible.
- With V4R4, TCP/IP can take advantage of larger buffers. Prior to V4R4, the **TCP/IP buffer size** (TCPRCVBUF and TCPSNDBUF on the CHGTCPA or CFGTCP command) was recommended to be increased from 8K bytes to 64K bytes to maximize data rates. When transferring large amounts of data with V4R4, you may receive higher throughput by increasing these buffer sizes up to 8MB.
- To receive the full benefit of the performance improvements in V4R4, it is essential that Ethernet is used and that the **TCPONLY parameter in the LIND have a value of *YES**. This allows the IOP to only have the TCP/IP version of its microcode active. This allows the IOP and device driver running on the AS/400 CPU to run in optimized mode. If other high-level protocols are active on that line, like APPC, then this parameter must be set to *NO for functional reasons.
- If Ethernet is used with TCPONLY(*NO) or if TRLAN is used, then you will realize only a portion of the V4R4 performance improvements in terms of CPU time reduction.
- The TCP/IP communications support also had significant potential capacity increase because of better scalability in V4R4 due to software contention reduction. NetPerf measurements executed in the Rochester performance lab on N-way processors indicate significant overall capacity increases over V4R3.
- The minimum Request/Response round trip delay is about 1 millisecond for TCP/IP using an Ethernet IOP (with TCPONLY=*YES). The CPW value for the CPU, the size of Request/Response, along with the load on the system will impact these times. These IOP delays are most noticeable in user transactions that contain many individual communications I/Os (like database serving). Having a fast IOP is critical to response time for these client/server environments.
- Connections and closes using sockets are significantly more expensive than normal sends and receives. Limit the number of times that new connections must be established. (i.e., leave the connection up if possible). Note from the data in Table 5.1 that a SSL Connect/Request/Response uses over 6 times more CPU than with a simple Request/Response when the connection is already in place.
- Application time for transfer environments, including accessing a data base file, decreases the maximum potential data rate. Because the CPU has additional work to process, a smaller percentage of the CPU is available to handle the transfer of data. Also, serialization from the application's use of both database and communications will reduce the transfer rates.
- For optimum performance with FTP, make sure the MTU size is as large as possible. ASCII transfers are slower because of character conversion from EBCDIC to ASCII prior to the data being sent. For

PUTs and GETs with FTP from IFS, the best possible transfer rate usually ranges from 20-25 Mbps.

• Communications performance improvements that were introduced in V3R1 for APPC, TCP/IP, UDP/IP, FTP, and TELNET are discussed in AS/400 Performance Capabilities Reference (V3R2), ZC41-8166.

Capacity Planning with TCP/IP:

The following table provides some rough capacity planning information for communications when using Sockets, TCP/IP, and Ethernet. Comparisons with SSL and VPN are also included. Note that it is always better to project the performance of an application from measurements based on that same application, but that is not always possible. This information is of similar type to that provided in Chapter 6, Web Serving Performance. There are also a couple of capacity planning examples in that chapter. Note that this is based on the NetPerf workload, which is a primitive-level workload. The application does nothing other than to issue sockets APIs. A real user application will have this type of processing only as a percentage of the overall workload.

Table 5.1. V4R4 AS/400 TCP/IP Capacity Planning						
	Capacity Metric					
Transaction Type:	(transactions / second / CPW)					
	Non-secure	SSL	VPN	VPN (ESP		
		(RC4 / MD5)	(AH with MD5)	with DES / MD5)		
Request/Response	22.6	12.2	6.5	3.8		
(RR) 1 Byte						
Request/Response	3.2	.4	.3	.1		
(RR) 16K Bytes						
Asym. Connect/Request/Response	3.0	.5	.6	.3		
(ACRR) 8K Bytes						
Stream	10.1	.9	.5	.2		
16K Bytes						

Notes:

• Based on measurements with the NetPerf workload using an AS/400 Model 170/2386 with V4R4

• The data in the table reflects AS/400 as a server (not a client)

• The data reflects Sockets, TCP/IP, and 100 Mbps Ethernet with TCPONLY(*YES) configured. Variation of the protocol or the TCPONLY parameter will prvide different performance.

• SSL measurements used 128-bit RC4 symmetric cypher and MD5 message digest with 1024-bit RSA public/private keys. CRR data assumes a regular SSL handshake, not a full SSL handshake.

• VPN measurements used transport mode, 56-bit DES symmetric cypher and MD5 message digest with manually keyed RSA public/private keys.

• CPW is the "Relative System Performance Metric" found in Appendix D "AS/400 CPW Values"

• This is only a rough indicator for capacity planning. CPU capacities do not scale exactly by CPW; therefore, actual results may differ significantly.

For example, if a customer has an application that uses SSL to establish a connection to an AS/400 server (model 170/2386), issue a request, receive an 8K byte response, and close the connection; and wishes to use about 20% of the overall CPU for the network processing portion, then note the following calculation: 460 CPW * 0.5 trans/sec/CPW * 20% = 46 transactions per second.

Table 5.2	V4R4 AS/400 SSL and VPN Relative CPU	Time
1 <i>ubie</i> 5.2.	V4R4 AS/400 SSL and VI N Relative CI U	1 11116

	Relative CPU Time (Scaled to Non-Secure for each transaction type)					
Transaction Type:						
	Non-secure	SSL	VPN	VPN (ESP		
		(RC4/MD5)	(AH with MD5)	with DES/MD5)		
Request/Response	1.0 w	1.8 w	3.5 w	6.0 w		
(RR) 1 Byte						
Request/Response	1.0 x	7.9 x	10.6 x	29.7 x		
(RR) 16K Bytes						
Asym. Connect/Request/Response	1.0 y	5.1 y	5.1 y	10.8 y		
(ACRR) 8K Bytes						
Stream	1.0 z	11.1 z	21.5 z	46.5 z		
16K Bytes						

Notes:

• Based on measurements with the NetPerf workload using an AS/400 Model 170/2386 with V4R4

- The data in the table reflects AS/400 as a server (not a client)
- The data reflects Sockets, TCP/IP, and 100 Mbps Ethernet with TCPONLY(*YES) configured. Variation of the protocol or the TCPONLY parameter will prvide different performance.
- SSL measurements used 128-bit RC4 symmetric cypher and MD5 message digest with 1024-bit RSA public/private keys.. CRR data assumes a regular SSL handshake, not a full SSL handshake.
- VPN measurements used transport mode, 56-bit DES symmetric cypher and MD5 message digest with manually keyed RSA public/private keys.

• CPW is the "Relative System Performance Metric" found in Chapter 2, "AS/400 System Capacities and CPW"

- This is only a rough indicator for capacity planning. CPU capacities do not scale exactly by CPW; therefore, actual results may differ significantly.
- w, x, y and z are scaling constants, one for each NetPerf workload type.

SSL and VPN Performance Information:

- From Table 5.2, note the CPU Time required to process transactions in a secure mode. Some overheads are fixed while some are size related. The fixed overheads include the handshakes needed to establish a secure SSL connection. The variable overhead is based on the number of bytes that need to be encrypted/decrypted, the size of the public key, the type of encryption, and the size of the symmetric key.
- Use SSL and Sockets APIs wisely to minimize the number of secure transactions for a given application. Clearly the secure transactions require significantly more CPU time and will reduce overall transaction capacity.
- When a client makes a secure connection with SSL for the first time, additional handshakes and certificate processing must occur. This is referred to as the *full SSL handshake*. Once this has been done, and this client's information can stay in the server's session key cache, then *regular SSL handshakes* occur. Table 5.1 reflects regular SSL handshakes for the Connect/Request/Response scenario. The full SSL handshake can consume about 20 times more CPU than the regular SSL handshake.
- The SSL information provided in Table 5.1 and Table 5.2 uses 128-bit encryption with RC4/MD5 and has a key ring file with a 1024-bit public key. Other cypher suites and public key sizes will perform differently.
- Client authentication requested by the server is quite expensive in terms of CPU and should only be requested when needed.

- VPN works at the IP layer rather than the socket layer as with SSL. Hence, it is typically used to secure a broader class of data than SSL all of the data flowng between two systems rather than, for example, just the data between two applications. Other important differences include SSL does not protect UDP data, SSL cannot automatically generate new encryption keys (dynamic VPN connection) and securing a connection using VPN is completely transparent to the application.
- Under the covers, VPN eventually uses the same BSAFE routines as SSL.
- The performance of VPN will vary according to the level of security applied. In general, configure the lowest level of security demanded by your application
- In many cases data only needs to be authenticated. While VPN-ESP can perform authentication, AH-only affects system performance just half as much as the ESP with athentication and encryption. Another advantage of using AH-only is that AH authenticates the entire datagram, ESP, on the other hand, does not authenticate the leading IP header or any other informatin that comes before the ESP header.Packets that fail authentication are discarded and are never delivered to upper layers. This greatly reduces the chances of successful denial of service attacks.
- The VPN information provided in Table 5.1 and Table 5.2 uses transport mode, 56-bit encryption with DES/MD5 and has a key ring file with manually public/private keys. Other cypher suites will perform differently.

5.2 APPC, ICF, CPI-C, and Anynet Performance Information

APPC, ICF, CPI-C, and Anynet:

- Many general performance tips listed in the TCP/IP section are also pertinant to APPC.
- Ensure that APPC is configured optimimally for best performance:

LANMAXOUT on the CTLD (for APPC environments): This parameter governs how often the sending system waits for an acknowledgement. Never allow LANACKFRQ on one system to have a greater value than LANMAXOUT on the other system. The parameter values of the sending system should match the values on the receiving system.

In general, a value of *CALC (i.e., LANMAXOUT=2) offers the best performance for interactive environments, and adequate performance for large transfer environments.

For large transfer environments, changing LANMAXOUT to 6 may provide a significant performance increase

LANWNWSTP for APPC on the controller description (CTLD): If there is network congestion or overruns to certain target system adapters, then increasing the value from the default=*NONE to 2 or something larger may improve performance.

MAXLENRU for APPC on the mode description (MODD): If a value of *CALC is selected for the maximum SNA request/response unit (RU) the system will select an efficient size that is compatible

with the frame size (on the LIND) that you choose. The newer LAN IOPs support IOP assist. Changing the RU size to a value other than *CALC may negate this performance feature.

- In general TCP/IP provides better performance with V4R4. Some APPC APIs provide blocking (e.g., ICF and CPI-C), therefore scenarios that include repetitive small puts (that may be blocked) may achieve much better performance.
- A large transfer with the server AS/400 system sending each record repetitively using the default blocking provided by OS/400 to the client AS/400 system provides the best level of performance.
- A large transfer with the server AS/400 system flushing the communications buffer after each record (FRCDTA keyword for ICF) to the client AS/400 system consumes more CPU time and reduces the potential data rate. That is, each record will be forced out of the server system to the client system without waiting to be blocked with any subsequent data. Note that ICF and CPI-C support blocking, Sockets does not.
- A large transfer with the server AS/400 system sending each record requiring a synchronous confirm (e.g., CONFIRM keyword for ICF) to the client AS/400 system uses even more CPU and places a high level of serialization reducing the data rate. That is, each record is forced out of the server system to the client system. The server system program then waits for the client system to respond with a confirm (acknowledgement). The server application cannot send the next record until the confirm has been received.
- Compression with APPC should be used with caution and only for slower speed WAN environments. Many suggest that compression should be used with speeds 19.2 kbps and slower and is dependent on the data being transmitted (# of blanks, # and type of repetitions, etc.). Compression is very CPU-intensive. For the CPB benchmark, compression increases the CPU time by up to 9 times. RLE compression uses less CPU time than LZ9 compression (MODD parameters).
- ICF and CPI-C have very similar performance for small data transfers.
- ICF allows for locate mode which means one less move of the data. This makes a significant difference when using larger records.
- The best case data rate is to use the normal blocking that OS/400 provides. For best performance, the use of the ICF keywords force data and confirm should be minimized. An application's use of these keywords has its place, but the trade-off with performance should be considered. Any deviation from using the normal blocking that OS/400 provides may cause additional trips through the communications software and hardware; therefore, it increases both the overall delay and the amount of resources consumed.
- Having ANYNET = *YES causes extra CPU processing. Only have it set to *YES if it is needed functionally; otherwise, leave it set to *NO.
- For send and receive pairs, the most efficient use of an interface is with it's "native" protocol stack. That is, ICF and CPI-C perform the best with APPC, and Sockets performs best with TCP/IP. There is CPU time overhead when the "cross over" is processed. Each interface/stack may perform differently depending on the scenario.

- Copyfile with DDM provides an efficient way to transfer files between AS/400s. DDM provides large blocking which limits the number of times the communications support is invoked. It also maximizes efficiencies with the data base by doing fewer larger I/Os. Generally, a higher data rate can be achieved with DDM compared with user-written APPC programs (doing data base accesses) or with ODF.
- When ODF is used with the SNDNETF command, it must first copy the data to the distribution queue on the sending system. This activity is highly CPU-intensive and takes a considerable amount of time. This time is dependent on the number and size of the records in the file. Sending an object to more than one target AS/400 only requires one copy to the distribution queue. Therefore, the realized data rate may appear higher for the subsequent transfers.
- FTS is a less efficient way to transfer data. However, it offers built in data compression for linespeeds less than a given threshold. In some configurations, it will compress data when using LAN; this significantly slows down LAN transfers.

5.3 LAN and WAN Performance Information

This section discusses the performance characteristics of local area network (LAN) protocols, lines, and IOPs.

LAN Media and IOP:

- No single station can or is expected to use the full bandwidth of the LAN media. It offers up to the media's rated speed of aggregate capacity for the attached stations to share. The CPU is usually the limiting resource. The data rate is governed primarily by the application efficiency attributes (for example, amount of disk accesses, amount of CPU processing of data, application blocking factors, etc.).
- LAN can achieve a significantly higher data rate than any other supported WAN protocol. This is due to the desirable combination of having a high media speed along with large frame sizes.
- When several sessions use a line or a LAN concurrently, the aggregate data rate may be higher. This is due to the inherent inefficiency of a single session in using the link.
- TCP/IP adaptive packet training was added beginning in V4R2. This enables the software to modify internal wait times to balance CPU utilization and round trip send/receive times. For example, the round trip time for a small LAN frame dropped from about 15 milli-seconds to less than 5 ms.
- In order to achieve good performance in a multi-user interactive LAN environment it is recommended to manage the number of active users so that LAN media utilization does not exceed 50% for TRLAN or 25% for Ethernet environments with multiple users because of media collisions resulting in thrashing. Operating at higher utilizations may cause poor response time due to excess queuing time for the line. In a large transfer environment where there is a small number of users contending for the line, at any given time a higher line utilization may still offer acceptable performance.
- There are several parameters in the line description and the controller description that play an important performance role.

- MAXFRAME on the line description (LIND) and the controller description (CTLD): Maximizing the frame size in a LAN environment is very important and supplies best performance for large transfers. Having configured a large frame size does not negatively impact performance for small transfers. Note that both the AS/400 system and the other link station must be configured for large frames. Otherwise, the smaller of the two maximum frame size values is used in transferring data. Bridges may also limit the maximum frame size. Note that the maximum frame size allowed is 16393 for TRLAN and that a smaller value is the default.
- TCPONLY on the line description (LIND): The parameter activates a higher-performance software feature which optimizes the way in which the IOP and the CPU pass data. This can be set to a value of *YES if TCP/IP is the only protocol to be used (e.g., not APPC).
- See the TCP/IP section for other related parameters.
- See the APPC section for other related parameters.
- When configuring an AS/400 system with communications lines and LANs it is important not to overload an IOP to avoid a possible system performance bottleneck.
- For interactive environments it is recommended not to exceed 60% utilization on a LAN IOP. Exceeding this threshold in a large transfer environment or with a small number of concurrent users may still offer acceptable performance. Use the AS/400 performance tools to measure utilization.
- Optimally configured, the 100 Mbps Ethernet IOP/IOA can have an aggregate transfer rate of up to 50 Mbps for TCPONLY(*NO) and up to 90 Mbps for TCPONLY(*YES). Multiple concurrent large transfers may be required to drive the IOP at that rate. (This assumes the use of the most recent IOP).
- Similarly in a web server environment using 100 Mbps Ethernet, the IOP capacity may be up to 120 hits/sec for TCPONLY(*NO) and 245 hits/sec for TCPONLY(*YES). This assumes nonsecure transactions and static pages of about 10K bytes each.
- The TRLAN IOP can support aggregate transfer rates of almost 16 Mbps, which is media speed. (This assumes the use of the most recent IOP).
- 100 Mbps Ethernet may provide a performance improvement over 10 Mbps Ethernet in terms of higher data rates for large transfer environments. It may also provide better overall performance within an establishment if the previous 10 Mbps Ethernet media was nearing its media utilization threshold.
- It is especially important to have a high-capacity IOP available for file serving, data base serving, web serving or for environments that have many communications I/Os per transaction. This characteristic will also minimize the overall response time.
- Higher-performing TRLAN IOP/IOAs have the potential to overrun lesser capacity TRLAN IOP/IOAs. Many re-transmissions and time-out conditions exist here. Check the AS/400 performance tools for these statistics. For APPC, this can be minimized or avoided by limiting the LANACKFRQ and LANMAXOUT parameters to 1 and 2, respectively, which are the default values.

- A given model of the AS/400 system can attach multiple IOPs up to a given maximum number. It is important to distribute the workload across several IOPs if the performance capability of a single IOP is exceeded. There are also some limitations on the number of stations that can be configured through a single LAN connection.
- The larger maximum frame size gives 16Mbit Token Ring emulation over ATM the advantage vs. Ethernet emulation over ATM.

WAN Line and IOP:

- Typically WAN refers to communications lines running at 64Kbps or slower. In recent years, other WAN types (like Frame Relay) have increased media speed up to several Mbps.
- In many cases, the communications line is the largest contributor to overall response time. Therefore, it is important to closely plan and manage its performance. In general, having the appropriate line speed is the most key consideration for having best performance.
- A common misconception exists in sizing systems with communications lines. It is incorrect to believe that each attached line consumes CPU resource in a uniform fashion, and therefore, exact statements can be made about the number of lines that any given AS/400 model can support. For example, if the sales pages say that a particular AS/400 model supports 64 lines, it does not mean that any given customer can run their workload fully utilizing those 64 lines. It is merely a rough guideline stating the suggested maximum for that model (in some cases, it is the maximum configuration possible).
- Communications applications consume CPU and IOP resource (to process data, to support disk I/O, etc.) and communications line resource (to send and receive data or display I/O). The amount of line resource that is consumed is proportional to the total number of bytes sent or received on the line. Some additional CPU resource is consumed to process the communications software to support the individual sends (puts or writes) and receives (gets or reads). Communications IOP resource is also consumed to support the line activity.

So the best question to ask is NOT "How many lines does my system support?", but rather, "How many lines does my workload require, and what AS/400 model is required to accommodate this load?".

- To estimate the utilization of a half duplex line: utilization = (bytes in + bytes out) * 800 / time / linespeed where time = total # of seconds and linespeed = the speed of the line in bits per second
- For a full duplex line (e.g., X.25, ISDN), the AS/400 Performance Tools report utilization as follows: Utilization = (bytes in + bytes out) * 400 / time / linespeed

For example, if the send direction is 100% busy and the receive direction is 0% busy, the Performance Tools will report an overall 50% line utilization.

• The system usually can drive the line to a high utilization for applications that transfer a large amount of data. The difference of the data rate and the line speed is due to the overhead of header bytes, line

turn around 'dead' time, and application serialization.

- When several sessions use a line concurrently, the aggregate data rate may be higher. This is due to the inherent inefficiency of a single session in using the link. In other words, when a single job is executing disk operations or doing non-overlapped CPU processing, the communications link is idle. If several sessions transfer concurrently, then the jobs may be more interleaved and make better use of the communications link.
- For interactive environments, keeping line utilization below 30% is recommended to maintain predictable and consistent response times. Exceeding 50% line utilization will usually cause unacceptable response times. The line utilization can be measured with the AS/400 performance tools.
- For large transfer environments, or for environments where only a small number of users are sharing a line, having a higher line utilization may yield acceptable response times. In fact, maximizing line utilization means maximizing throughput for that single job.
- For large transfers, use large frame sizes for best performance. Fewer frames make more efficient use of the CPU, the IOP, and the communications line (higher effective data rate).
- To take advantage of these large frame sizes, they must be configured correctly. The MAXFRAME parameter on the LIND must reflect the maximum value. For X.25, the DFTPKTSIZE and MAXFRAME must be increased to its maximum value. Also, go to the APPC and TCP sections to ensure other related parameters are optimized.
- Configuring a WAN line as full-duplex may provide a higher throughput for certain applications that can take advantage of that, or for multiple-user scenarios.
- In general, the physical interface does not noticeably affect performance for a given protocol assuming that all other factors are held constant (e.g., equal line speeds). For example, if SDLC is used with a line speed of 19.2 kbps, it would not matter if a V.35, RS232, or an X.21 interface was used (all other factors held constant).
- For SDLC environments, polling is an important consideration. Parameters can be adjusted to change the rate at which a line is polled. Polls consist of small frames sent across the line and are processed by the IOPs. Therefore, polling contributes to line utilization and IOP utilization.
- The CPU usage (i.e., CPU time per unit of data) for SDLC and X.25 is similar. Depending on the application design, BSC and Async may require more CPU.
- The CPU usage for high speed WAN connections is similar to "slower speed" lines running the same type of work. As the speed of a line increases from a traditional low speed to a high speed (e.g., 1-2 Mbps), performance characteristics may change.
 - ✤ Interactive transactions may be slightly faster
 - ✤ Large transfers may be significantly faster
 - ✤ A single job may be too serialized to utilize the entire bandwidth
 - ✤ High throughput is more sensitive to frame size
 - High throughput is more sensitive to application efficiency
 - System utilization from other work has more impact on throughput

- The WAN-capable IOPs handle the load with a relatively low IOP utilization and generally won't be the system performance capacity bottleneck.. However, you may check the IOP's utilization by using the Performance Monitor.
- For interactive environments it is recommended not to exceed 60% utilization on the communications IOP. Exceeding this threshold in a large transfer environment or with a small number of concurrent users may still offer acceptable performance. Use the AS/400 performance tools to measure utilization.
- Even though an IOP can support certain configurations, a given AS/400 model may not have enough system resource (for example, CPU processing capacity) to support the workload over the lines.
- In communications environments where errors are common, the use of smaller frame sizes may offer better performance by limiting the size of the re-transmissions. Having errors may also impact the number of communications lines that can run concurrently.
- The values for IOP utilization in SDLC environments do not necessarily increase consistently with the number of work stations or with the amount of workload. This is because an IOP can spend more time polling when the application is not using the line. Therefore, it is possible to see a relatively high IOP utilization at low throughput levels.

5.4 Work Station Connectivity Performance Information

There are many ways to attach work stations (WS) to the AS/400 via communications. Each type can have different overheads with the CPU or the media and can have other unique performance characteristics.

Work Station Connectivity:

- Interactive transactions include CPU processing for WS connection and application processing. If the application is "light" (similar to a commercial workload), then the performance impact of how WS are connected can be significant. Here, the percentage of CPU consumed to process screen I/O is greater. If the application is "complex" (significantly more CPU processing per transaction), then the performance impact of connectivity type is less significant and the percentage of CPU consumed to process screen I/O is less.
- Attaching WS through communications consumes more CPU than 5250 local WS support does. Keep in mind that the actual overhead may also vary significantly with changes in the data stream (number of I/Os, number of bytes, number of fields, and other screen I/O characteristics). For the following comparisons, we can compare the overall amount of CPU processing done for a "light" application (i.e., the amount to handle the screen I/O plus that to process the application). These comparisons assume that only this application is running on the system without any other workload consuming CPU. All comparisons are done with respect to the 5250 local WS baseline.
 - ✤ WS attached with 5250 target-side DSPT, with remote work stations with the 5495 controller, or with CA/400 increase overall application CPU requirements for communications by about 10% and therefore reduce potential AS/400 capacity by 10%.

- WS attached with TELNET increase overall application CPU requirements by about 25% and therefore reduce potential AS/400 capacity by 20%. This is due to additional CPU processing per transaction for TCP/IP software and TELNET. Note that using the IBM Network Station as a work station uses this method of attachment. Note that this overhead varies greatly based on the screen characteristics. This impact can vary greatly due to the characteristics of the data stream (# bytes, # fields, etc.).
- ✤ WS attached with VT100/VT220 increase overall application CPU requirements by about 60% and therefore reduce potential AS/400 capacity by about 40%. This is due to additional CPU processing for TCP/IP, TELNET, and data stream translation. This impact can vary greatly due to the characteristics of the data stream (# bytes, # fields, function keys, etc.).
- WS attached with 3270 Remote Attach, DHCF, NRF, or SPLS increase overall application CPU requirements by about 25% and therefore reduce potential capacity AS/400 capacity by 20%. This is due to additional CPU processing for communications and data stream translation per transaction. This impact can vary greatly due to the characteristics of the data stream (# bytes, # fields, complexity, etc.).
- Web server based packages that allow 5250/HTML sessions significantly increase CPU requirements by several fold and therefore reduce potential AS/400 capacity by several times. This is due to additional CPU processing for communications, the web server, and 5250 to HTML conversions.
- Passing Through an AS/400 to an application on another system has several possibilities:
 - 5250 DSPT (source side) is the baseline here. This "front-end" system only has to support the WS attachment processing and communications support to the server system with no application processing being it is just a source side.
 - TELNET (source side) uses several times more CPU time (2-5 times more depending on the screen characteristics) than 5250 DSPT because of additional processing for TCP/IP communications and more processing in the TELNET application.
 - ✤ APPN intermediate node routing (intermediate system passing transaction from source to target) has a similar CPU time to 5250 DSPT (source).
- **Twinaxial controllers provide better performance than ASCII controllers.** This is primarily due to the increased line speed. The conversion from ASCII to EBCDIC is performed in the ASCII controller, so it is not an impact to AS/400 CPU time. ASCII response time should be similar to the response time for a remote work station configuration with a similar line speed.
- Keep the line utilization below 30% for best performance when interactive users are attached. This will maintain predictable and consistent response times. Exceeding 50-60% line utilization will usually cause unacceptable response times.
- **Mixed interactive users and batch**: When interactive users and large transfers are running on a communications line concurrently, consider the following to keep interactive performance acceptable:

- 1. Use APPN transmission priority to prioritize the interactive users' transfers over that of the large transfer. (this is the preferred choice, as it does not penalize the large transfer when there is no interactive traffic)
- 2. Change the RU size to a lower value for the large transfer. This optimizes interactive response time at the expense of large transfer performance (note that overall CPU time will increase also for the large transfer).
- 3. Reducing the pacing values for the large transfer will also slow it down, allowing the interactive users more windows for getting on the line.

5.5 Opti-Connect for OS/400 Performance Information

OptiConnect/400 (OC/400) enables applications to transparently access databases on remote AS/400 systems. It allows applications written to access databases locally, to access them remotely with simple changes to file descriptions and no changes to the applications.

OC/400 provides an optimized DDM solution capable of efficient, low-latency, high bandwidth communication between AS/400 systems. OC/400 utilizes a fiber optic connected shared I/O bus and a SLIC shared bus device driver to transport data and messages. The efficiency of OC/400 allows horizontal growth solutions - increasing total capacity to shared data beyond the bounds of a single system. Prior to OC/400, DDM using ICF/APPC for transport has provided the remote data access function but the overhead of communication between systems was too high to achieve horizontal growth under heavy workloads. This overhead kept DDM from being a viable solution for user's horizontal growth. OC/400 offers customers this horizontal growth by enabling 2 to 32 AS/400 systems to share an I/O bus. Any or all systems can act as both an application machine (workstation server) for accessing another systems data or a database server providing data to other systems.

The following table shows system capacity growth potential when OC/400 is used in system environments with various levels of data base activity. This table only applies to the work that is being considered for OC/400 usage. The database server machine is assumed to have similar configuration like the single system has. These are only guidelines and results will vary depending on system environments. These guidelines were based on data from a previous release, but it is assumed that the relative ratios would still hold. The growth factors will also vary with the workload (read/write ratios, open/close frequency, etc.).

Table 5.3 OC/400 Horizontal Growth Potential							
OC/400 Horizontal Growth Potential							
Single Database Server and Multiple Application Machines (1 to 7)							
DB %	1-A	2-A	3-A	4-A	5-A	6-A	7-A
10	1.81	2.63	3.44	4.26	5.07	5.89	6.70
20	1.66	2.31	2.97	3.62	4.17		
30	1.52	2.03	2.55	2.78			
40	1.39	1.79	2.08				
50	1.29	1.57	1.67				
60	1.19	1.38	1.39				

Note:

• DB% = The percent CPU time allocated to database activity in a single system environment.

• 2-A = 3 systems composed of 2 application machines and 1 database server.

• Client Factor = 1.8 for this workload (ie, the OC/400 CPU associated with the client; eg, offloading 10% database would be replaced by 18% OC/400 CPU).

• Server Factor = 1.2 for this workload (ie, the OC/400 and related database CPU associated with the server; eg, offloading 10% database from the client would represent 12% CPU on the server).

• A typical customer environment includes 1 database server and 2 or more application machines.

• The columns show the capacity gain over a single system. For example, 20% DB with 2-A shows a value of 2.31. This value is the capacity growth achieved by using 3 systems; 1 as server and 2 as application machines over a single system. Any remaining capacity on the server not directed at the client(s) can do additional single system workload.

• All systems are asumed to have the same CPW value for this exercise.

1. Each time that a 'put' or 'get' is issued, OC/400 uses the shared I/O bus to exchange data with the remote system. DDM uses a communication link for accessing data on remote system.

- 2. V4R1 OC/400 performance should be equal or better than V3R6 (and also V3R1) when the capacity is scaled to CPW of the CPU.
- 3. OC/400 is a more efficient way than DDM (with communications) to access data on another (near by) AS/400.
- 4. At higher levels of throughput the system overhead for OC/400 increases in a non-linear fashion. This is due to multiprocessing contention when CPU utilization is higher. It is recommended that CPU utilization should not exceed 90% on either the application machine or database server system for this reason.
- 5. There is an overhead (called client/server factor) associated with using OC/400 both on the application and data base machine. This overhead varies and is dependent on factors like:
 - Percent of data base activity
 - Number of logical IOs per transaction
 - Usage of Journal and Commitment Control
 - Number of data base Open/Close
- 6. The Maximum length for the OC/400 fiber optic cables to the shared bus is 2KM, requiring all OC/400 systems to be within 2KM of the system providing the shared bus. Due to fiber latency of 5 nano-seconds per meter, it is recommended that systems should be within 300 meters of each other. This may avoid any performance impact due to line latency. DDM using APPC over LANs or WANs can provide greater distances.
- 7. The OC/400 adapter's capacity may be exceeded with the new V4R1 CPU models. You may use WRKOPCACT command or the PFRMON (logical I/O field) to observe the OC/400 ops/sec rate. From measurements with a workload that did small read operations (data not provided here), the following maximums were observed:
 - Features 2682/2685/2688 (20M cable): 10,200 ops/sec (1.8 Mbps)
 - Features 2682/2685/2688 (500M cable): 7,100 ops/sec (1.2 Mbps)
 - Features 2680/2683/2686 (2KM cable): 2,700 ops/sec (0.5 Mbps)

On an 8-way model 650 with the workload mentioned above, two OC/400 adapters were needed to fully drive the 8-way to full capacity. The PFRMON does not show the utilization of this adapter in the reports. So you should consider these maximums to see if the adapter is possibly the bottleneck. If so, adding additional adapters and cabling is possible.

The 2682/2685/2688 features provide a higher level of performance because of being associated with cables of 500M or less.

The length of the cable also impacts OC/400 capacity. Note from above that throughput drops with cable length by up to 30% going from the short to long cable with the 2682/2685/2688 features.

8. The OC/400 horizontal growth is achieved across multiple systems with some or all systems acting as servers. Any single database network of logical and physical files must reside on a single system, therefore multiple servers require multiple, separate database networks. The maximum

load supported across multiple systems is typically limited by the capacity of a database server - additional application systems can be added until the database server saturates.

- 9. When analyzing batch workloads, more attention must be paid to logical database I/O's, blocking, and pool size. Performance characteristics may be different for batch versus interactive use of OC/400. In particular, batch jobs are often very database intensive and accumulated CPU and delays for each logical I/O which must traverse the OC/400 shared bus can add significant CPU utilization and elapsed time to batch runs.
- 10. A detailed analysis to determine where to place the data files is very critical in batch workloads or read only files.
- 11. An OC/400 shared bus adapter can support up to 40 MBytes/second data rate. Measured bus/adapter utilization in majority of OC/400 testing has shown CPU or adapter limits are reached prior to reaching bus/adapter limits
- 12. In most distributed environments, including OC/400, there is a penalty for remote access; obviously this penalty is most severe when 100% of accesses are remote and may become negligible for lower percentages (eg. 10%). Therefore an optimum environment should minimize remote access by distributing users or data across systems in a manner to maximize local access. Replication of read-only files to application machines, distribution of data across multiple database servers, and location of batch work on application machines or database servers must be carefully considered when defining an OC/400 environment to optimize total system capacity.

5.6 NetPerf Workload Description

The NetPerf workload is a primitive-level function workload to explore communications performance. The NetPerf workload consists of C programs that run between a client AS/400 and a server AS/400. Multiple instances NetPerf can be executed over multiple connections to increase the system load. The programs communicate with each other using sockets or SSL programming APIs.

Whereas most 'real' application programs will process data in some fashion, these benchmarks merely copy and transfer the data from memory. Therefore, additional consideration must be given to account for other normal application processing costs (for example, higher CPU utilization and higher response times due to database accesses).

To demonstrate communications performance in various different ways, several scenarios with NetPerf are analyzed. Each of these scenarios may be executed with regular non-secure sockets or with secure SSL:

- 1. **Request/Response (RR)**: the client and server send a specificed amout of data back and forth over a connection that remains active. This is similar to client/server application environments.
- 2. Asymmetric Connect/Request/Response (ACRR): the client establishes a connection with the server, a single small request is sent to the server, and a response (of a specified size) is sent by the server, and the connection is closed. This is a web-like transaction.
- 3. **Stream (large data transfer):** the client repetitively sends a given amount of data to the server over a connection that remains active.

Chapter 6. Web Serving Performance

Performance information for web serving on the AS/400 and various types of web server transactions will be discussed in this section.

There are many factors that can impact overall performance (e.g., end-user response time, throughput) in the complex Internet environment, some of which are listed below:

- Web Browser
 - processing speed of the client system
 - ✤ performance characteristics of the Web browser
 - client application performance characteristics
- Communications network
 - speed of the communications links
 - capacity of any proxy servers
 - ✤ congestion of network resources
- AS/400 Web server
 - ✤ AS/400 processor speed
 - utilization of key AS/400 resources (CPU, IOP, memory, disk)
 - ✤ Web server performance characteristics
 - ✤ application (e.g., servlet) performance characteristics

The primary focus of this section will be to discuss the performance characteristics of the AS/400 as a server in a Web serving environment, providing capacity planning information and recommendations for best performance. Please refer to Chapter 5, "Communications Performance" Ofor related information.

Data accesses across the Internet differ distinctly from accesses across 'traditional' communications networks. The additional resources to support Internet transactions by the CPU, IOP, and line are significant and must be considered in capacity planning. Typically, in a traditional network:

- there is a request and response (between client and server)
- connections/sessions are maintained between transactions
- networks are tuned to use large frames

For web transactions, there are a dozen or more line transmissions (including acknowledgements) per transaction:

- a connection is established/closed for each transaction
- there is a request and response (between client and server)
- networks typically have small frame (MTU) sizes
- one user transaction may contain separate internet transactions
- secure transactions are frequent and consume more resource

The information that follows is based on performance measurements and analysis done in the internal IBM performance lab. The raw data is not provided here, but the highlights, general conclusions, and recommendations are included. Results listed here do not represent any particular customer environment. Actual performance may vary significantly from what is provided here. Note that these workloads, along with other published benchmark data (from other sources) are always measured in best-case environments (e.g., local LAN, large MTU sizes). Real internet networks typically have higher contention, MTU size limitations, and intermediate network servers (e.g., proxy, SOCKS).

6.1 Web Serving with the HTTP Server

The Hypertext Transfer Protocol (HTTP server) allows AS/400 systems attached to a TCP/IP network to provide objects to any Web browser. At a high level, the connection is made, the request is received and processed, the data is sent to the browser, and the connection is ended. The HTTP server jobs and the communications router tasks are the primary jobs/tasks involved (there is not a separate user job for each attached user).

Workload Description and Data Interpretation: The workload is a program that runs on a client work station. The program simulates multiple Web browser clients and repetitively issues 'URL requests' to the AS/400 Web server. The number of simulated clients can be adjusted to vary the offered load. Each of the transaction types listed in the tables serve about 1000 bytes:

- StaticPage: serves a static page via the HTTP server. This information can be accessed from the web server's cache of specified IFS files.
- CGI (HTML): invokes a CGI program that accesses data from IFS and serves a simple HTML page via the HTTP server. This runs in a named activation group.
- CGI (SQL): invokes a CGI program that performs a simple SQL request and serves the result via the HTTP server. This runs in a named activation group.
- Persistent CGI: invokes a CGI program that receives a handle supplied by the browser, accesses data from IFS, and serves a simple HTML page via the HTTP server.
- Net.Data (HTML): invokes the Net.Data program that serves a simple HTML page via the HTTP server.
- Net.Data (SQL): invokes the Net.Data program that performs a simple SQL request and serves the result via the HTTP server.
- Servlet: invokes a Java servlet that accesses data from IFS and serves a simple HTML page via the HTTP server.

Each of the above can be served in secure or non-secure fashion. "*Relative CPU time*" is the average AS/400 CPU time to process the transaction for each specific scenario. "*AS/400 Capacity (hits/sec/CPW)*" is the capacity metric used to estimate the capacity of any AS/400 model . Note that transactions/sec/CPW can be used interchangeably with hits/sec/CPW. An example exists in the conclusions. "*Secure:Nonsecure CPU time ratio*" indicates the extra CPU processing required to execute a given

transaction in a secure mode.

The CGI programs were compiled using a "named" activation group. For more information on program activation groups refer to *AS/400 ILE Concepts*, *SC41-5606*.

Web Serving Performance Measurements: The following tables provide a summary of the measured performance data. These charts should be used in conjunction with the rest of the information in this section for correct interpretation. Results listed here do not represent any particular customer environment. Actual performance may vary significantly from what is provided here.

Table 6.1. V4R4 AS/400 Web Serving Capacity Planning							
	Nonsecure		Secure				
Transaction Type:	Capacity Metric: hits/sec/CPW	Relative CPU Time	Capacity Metric: hits/sec/CPW	Secure:Nonsecure CPU Time Ratio			
Static Page (cached)	1.86	0.6 x	0.58	3.2			
Static Page (not cached)	1.18	1.0 x	0.48	2.5			
CGI (HTML)	0.44	2.7 x	0.28	1.6			
CGI (SQL)	0.43	2.7 x	0.28	1.5			
Persistent CGI	0.44	2.7 x	0.25	1.8			
Net.Data (HTML)	0.24	4.9 x	0.19	1.3			
Net.Data (SQL)	0.15	7.9 x	0.13	1.2			
Servlet	0.40	2.9 x	0.28	1.4			

Note:

• IBM HTTP Server for AS/400; V4R4; 100Mbps Ethernet; with TCPONLY(*YES)

• Based on measurements from an AS/400 Model 720-2062

• Static page caching done with IBM HTTP server (WRKHTTPCFG)

All requests cached for Net.Commerce

• 1KB data is served for each of the transaction types

• Data assumes no access logging

• Static page serving is served from the root directory of IFS

• CGI programs compiled with "named" activation group

• Secure measurements done using Secure Sockets Layer (SSL) with 40-bit RC4 encryption

• transactions/second/CPW can be used interchangeably with hits/sec/CPW

• CPW is the "Relative System Performance Metric" found in Chapter 2, "AS/400 System Capacities and CPW"

• Web server capacities may not scale exactly by CPW, therefore, results may differ significantly from those listed here

• NA = not available



Figure 6.1 AS/400 Web Serving V4R4 Relative Capacities

Web Serving Performance Tips and Techniques:

1. **V4R4** provides a performance improvement of up to 70% over that of V4R3 (with similar hardware). This is mostly due to improvements in the IBM HTTP Server and TCP/IP performance. For static pages that are not cached, V4R4 provides up to 7% more capacity. For static pages that are cached, V4R4 provides up to 20% more capacity. For CGI and Net.Data transactions, V4R4 provides up to 70% more capacity.

V4R3 provided a performance improvement in capacity of up to 65% over that of V4R2 (with similar hardware). This is mostly due to the improved efficiency of the IBM HTTP Server over that of the ICS/400 from V4R2. For static pages that are not cached, V4R3 provides up to 20% more capacity. For static pages that are cached, V4R3 provides up to 65% more capacity. There were also significant performance improvements for Net.data and CGI with named activations in V4R3.

2. **Web Serving Capacity (Example Calculations):** Throughput for web serving is typically discussed in terms of the number of hits/second or transactions/second. Typically, the CPU will be

the resource that determines overall system capacity. If the IOPs become the resource that limits system throughput, then the number of IOPs supporting the load could be increased. For system configurations where the CPU is the limiting resource, Table 6.1 above can be used for capacity planning. Use these high-level estimates with caution. They do not take the place of a complete capacity planning session with actual measurements of your particular environment. Remember that these example transactions are fairly trivial. Actual customer transactions may be significantly more complex and therefore consume additional CPU resources. Scaling issues for the server, the application, and the database also may come into consideration when using N-way processors with higher projected capacities.

Example 1: Estimating the capacity for a given model and transaction type: Estimate the system capacity by multiplying the CPW (relative system performance metric) for the AS/400 model with the appropriate hits/second/CPW value (the capacity metric provided in the table).

Capacity = *CPW* * *hits/sec/CPW*.

For example, a 170-2386 rated at 460 CPW doing web serving with CGI programs, would have a capacity of 202 trans/sec (460 x 0.44 = 202). This assumes that the entire capacity of the system would be allocated to Web serving. If other work will also be on the system, you must pro-rate the CPU allocation. For example, if only 25% of the CPU is allocated for Web serving, then it would have a web serving throughput of 50 trans/sec (460 x $0.44 \times 25\% = 50$).

Example 2: Estimating how many CPWs are required for a given web transaction load:

Characterize the transaction make-up of the estimated workload and the required transaction rate (in transactions/second). Estimate the CPWs required to support a given load by dividing the required transaction rate by the appropriate hits/second/CPW value (the capacity metric provided in the table).

Required CPWs = transaction rate / hits/sec/CPW.

For example, in order to support 175 CGI trans/sec, 398 CPWs would be required (175 / 0.44 = 398 CPWs). If a mixed load is being assessed, then calculate the required CPWs for each of the components and add them up. Select an AS/400 model that fits and allows enough room for future growth.

3. Net.Data:

- Net.Data is more disk I/O intensive than typical HTTP transactions. Therefore, more HTTP server jobs may be needed to provide the optimal level of system throughput.
- A Net.Data SQL macro is slower than an SQL CGI.bin. This is because the Net.Data SQL macro is interpreted while the SQL CGI.bin is compiled code. There are functional advantages in using an SQL macro.
 - direct reuse of existing SQL statements (no programming required)
 - provides the built-in ability to format SQL results
 - provides the ability to store SQL results in a table and pass the results to a different language environment (e.g., REXX).
- 4. **CGI and Persistent CGI**: Significant (perhaps as much as 6x) performance benefits can be realized by compiling into a "named" versus a "new" activation group. It is essential for good performance that CGI-based applications use named activation groups. Refer to the AS/400 ILE Concepts for more details on activation groups.

Persistent CGI is specific to applications needing to keep state information across web transactions. Don't confuse persistent CGI with a way to improve the performance of your CGI program. You'll notice in the earlier table that the performance of CGI is nearly identical to that of persistent CGI due to the advantages gained by running in a "named" activation group.

- 5. **Web Server Cache for IFS Files**: Serving static pages that are cached can increase web server capacity by about 50%. Ensure that highly used files are selected to be in the cache (WRKHTTPCFG).
- 6. **Page size**: The data in the tables assumes about 1K bytes being served. If the pages are larger, more bytes are processed, CPU processing per transaction significantly increases, and therefore the transaction capacity metrics would be reduced.
- 7. **Response Time (general):** User response time is made up of Web browser (client work station) time, network time, and server time. A problem in any one of these areas may cause a significant performance problem for an end-user. To an end-user, it may seem apparent that any performance problem would be attributable to the server, even though the problem may lie elsewhere.

It is common for pages that are being served to have imbedded images (e.g., gifs). Each of these separate Internet transactions adds to the response time since they are treated as independent HTTP requests and can be retrieved from various servers (some browsers can retrieve multiple URLs concurrently).

- 8. HTTP and TCP/IP Configuration Tips:
 - a. The **number of HTTP server jobs**: The CHGHTTPA command has parameters that specify the minimum and maximum number of server jobs. This is a system-wide value. The WRKHTTPCFG command also can specify similar values (MaxActiveThreads and MinActiveThreads). These values would override the values that are set via CHGHTTPA and would be for a given configuration. The reason for having multiple server jobs is that when one server is waiting for a disk or communications I/O to complete, a different server job can process another user's request. Also, for N-way systems, each CPU may simultaneously process server jobs. The system will adjust the number of servers that are needed automatically (within the bounds of the minimum and maximum parameters).

The values specified are the number of "child or worker" threads. Typically, 5 server threads are adequate for smaller systems (100 CPWs or less). For larger systems dedicated to HTTP serving, increasing the number of servers to 10 or more may provide better performance. A starting point for the maximum number of threads can be the CPW value divided by 20. Try not to have more than what is needed as this may cause unnecessary system activity.

b. The **maximum frame size parameter** (MAXFRAME on LIND) can be increased from 1994 bytes for TRLAN (or other values for other protocols) to its maximum of 16393 to allow for larger transmissions. Typically documents are larger than 1994 bytes.

c. The **maximum transmission unit (MTU) size** parameter (CFGTCP command) for both the route and interface affect the actual size of the line flows. Increasing these values from 576 bytes to a larger size (up to 16388) will most likely reduce the overall number of transmissions, and therefore, increase the potential capacity of the CPU and the IOP.

Similar parameters also exist on the Web browser. The negotiated value will be the minimum of the server and browser (and perhaps any bridges/routers), so increase them all.

- d. Increasing the **TCP/IP buffer size** (TCPRCVBUF and TCPSNDBUF on the CHGTCPA or CFGTCP command) from 8K bytes to 64K bytes may increase the performance when sending larger amounts of data. If data coming into the server is simply requests, increasing TCPRCVBUF may not provide any benefit.
- e. **Secure Web Serving**: Secure web serving involves additional overhead to the server. Additional line flows occur (fixed overhead) and the data is encrypted (variable overhead proportional to the number of bytes). Note the capacity factors in the tables above comparing non-secure and secure serving. For simple transactions (e.g., static page serving) the impact of secure serving is 2x or more based on the number of bytes served. For complex transactions (i.e., CGI or Net.Data), the overhead is in the range of 15-40%.
- f. E-Business applications typically yield a variety of complex transactions. These transactions have sub-transactions made up of static pages, CGI, Net.Data, etc. Capacity planning for these is more complex and warrants a careful analysis of the make-up of the transactions. The data from the tables can assist with this analysis.
- g. **Error and Access Logging**: Having logging turned on causes a small amount of system overhead (CPU time, extra I/O). Turn logging off for best capacity. Use the WRKHTTPCFG command to make these changes.
- h. **Name Server Accesses**: For each Internet transaction, the server accesses the name server for information (IP address and name translations). These accesses cause significant overhead (CPU time, comm I/O) and greatly reduce system capacity. These accesses can be be eliminated by using the WRKHTTPCFG command and adding the line "DNSLookUp Off".
- 9. **HTTP Server Memory Requirements**: Follow the faulting threshold guidelines suggested in the work management guide by observing/adjusting the memory in both the machine pool and the pool that the HTTP servers run in (WRKSYSSTS).

Factors that may significantly affect the memory requirements include using larger document sizes, using CGI.bin programs, and using Net.Data.

10. **AS/400 Model Selection**: Use the information provided in this section along with the characterization of your HTTP workload environment in a capacity planning exercise (perhaps with BEST/1) to choose the appropriate AS/400 model. All the tasks, jobs, and threads associated with HTTP serving are 'non-interactive', so AS/400e servers or AS/400 Advanced Servers would provide the best price/performance (unless other interactive work is present on the system).

11. **File System Considerations**: Web serving performance varies significantly based on which file system is used. Each file system has different overheads and performance characteristics. Note that serving from the ROOT or QOPENSYS directories provide the best system capacity. If Web page development is done from another directory, consider copying the data to a higher-performing file system for production use.

The web serving performance of the non-thread-safe file systems is significantly less than the root directory. Using QDLS or QSYS may decrease capacity by 2-5 times. For a more detailed discussion of IFS performance, please refer to the IFS section from the V3R2 version of this document.

- 12. **File Size Considerations**: The connect and disconnect costs are similar regardless of size, but cost for the transmission of the data with TCP/IP and the IFS access vary with size. As file size increases, the IOP is more efficient by being able to achieve a higher aggregate data rate. However, being larger, the files require more data frames, thus causing the hits/sec capacity for the IOP to go down accordingly.
- 13. **Communications/LAN IOPs**: Since there are a dozen or more line flows per transaction, the Web serving environment utilizes the IOP more than other communications environments. Use the performance monitor (STRPFRMON) and the component report (PRTCPTRPT) to measure IOP utilization. Attempt to keep the average IOP utilization at 60% or less for best performance.

IOP capacity depends on file size and MTU size (make sure you increase the maximum MTU size parameter). Additional information on communications/LAN IOP performance can be found in section LAN.

The 2619 or the 2617 LAN IOPs have a capacity of roughly 70 hits/sec when serving small (e.g., 1K byte) nonsecure pages (keep in mind that each hit contains a dozen or so line flows). Using Ethernet or TRLAN IOPs from V4R1 or more recent, have capacities in the 100-130 hits/sec range. If 100Mb Ethernet is used and the TCPONLY parameter in the LIND has a value of *YES, then capacities of up to 250 hits/sec may be seen.

On larger AS/400 models, the comm/LAN IOP may become the bottleneck before the CPU does. If additional HTTP capacity is needed, multiple IOPs (with unique IP addresses) could be configured. The overall workload would have to be 'manually' balanced by Web browsers requesting documents from a set of interfaces. The load can also be balanced across multiple IP addresses by using DNS (domain name server).

6.2 Net.Commerce Performance

Use the IBM AS/400 Workload Estimator to predict the capacity characteristics for Net.Commerce performance. Work with your marketing representative to utilize the tool at: *http://techsupport.rchland.ibm.com/supporthome.nsf/document/16533356*.

6.3 Firewall Performance

Using the Integrated PC Server as a firewall provides additional value-add for the AS/400 as a web server. The firewall can be on the same system as the web server or on a different system within the network. With the Integrated PC Server handling the firewall activity, the AS/400 CPU is not significantly impacted.

''Web server behind the firewall'': In this scenario the Integrated PC Server is performing packet filtering and allows HTTP traffic only through to the web server (also on the same AS/400). For an optimally configured system, having the firewall function active under a load, only slightly degrades the overall AS/400 web server capacity (compared with a similar, non-firewall configuration).

If a system is not optimally configured the decrease can be more significant. For example, if the MTU size is reduced to 500 bytes, then the impact of the firewall can be a 50% capacity reduction.

Chapter 7. Java Performance

Highlights:

- Introduction
- Improved Computational Performance (software)
- Reduced main storage footprint
- Unmatched Scalability
- Comparison to Existing Languages
- Java Performance -- Tips and Techniques
- Recommended Models and Capacity Planning

In traditional AS/400 applications, the performance of the application program itself is often a small contributor to overall performance. A large percentage of the execution is system services (e.g. Data Base Get Records) used by the application. Two ways to improve application performance are: 1) IBM improving OS/400, 2) The customer improving how the application uses the system services (especially data base) in OS/400.

Java, as part of its portability story, will have a higher percentage of the application's execution in Java programs and use less of a given Operating Service's function. Accordingly, the computational performance of the Java language is more important. The significance here is that OS/400's Java Virtual Machine (and the Direct Execution facility that translates Java bytecodes into directly executable programs) will be more important than in other languages. It will also be true that inefficient coding practices in the application and any third party Java code will also be more evident compared to traditional applications.

We'll discuss both computational performance (which is often measured by noncommercial programs) and commercial performance (Java implementing things recognizably commercial).

Last year at this time, there were no rigorous public benchmarks. For commercial Java processing, our soon-to-be-released jBOB benchmark partially fills the need. Its results will be supplemented with internal results. While Java benchmarks are emerging (some to great fanfare) their value in terms of predicting server performance of Java remains poor. Too many are strictly computational or else the Java contribution is minimal compared to nonJava path length.

7.1 Improved Computational Performance





Software Improvements

AS/400's Java product features improved computational performance over V4R3. Since V4R3 was also a substantial improvement over V4R2, Java performance has made great strides in a short period of time. The amount of improvement varies by the application.

However, a typical Java application, with minimal data base interactions, will improve approximately thirty to fifty percent over V4R3, on the same hardware.

This improvement comes about because of improvements in both code generation and in the allocation and garbage collection strategy of AS/400 Java's virtual machine.

Hardware Improvements

While V4R4 introduced a new series of models, the essential feature of these models is that they "filled in the gaps" from V4R3. In other words, while V4R3 concentrated on upgrades at the high end and the low end of the product line, using the new Northstar processors, V4R4 extended the range of the Northstar processors to the rest of the processor family. This will result in significant price-performance gains in the middle of the product line.

Of particular interest to Java customers is the new 2388. This introduces 2-way processing to the 170 model line for the first time. Since even a single-job Java program will have an independent, asynchronous garbage collection thread, two processors in a Java context might disproportionately benefit certain kinds of applications.

7.2 Reduced Main Storage Footprint

Given the current state of the art, accurate calibration of the main storage footprint on any machine is difficult. But, customers who experiment with V4R4 applications' performance versus their V4R3 applications should frequently see improved performance in smaller storage pools in most cases. Because it is possible to increase processor consumption to reduce main storage, and vice-versa, and because that tradeoff is application dependent, exact predictions cannot be made. Customers might try experiments where they reduce the pool size bit by bit, and observe the effects on Non Data Base Paging on the WRKSYSSTS display. (This sort of investigation should not be done on a mission critical machine). In many cases, the reduction in pool size should have little effect for a while, until some application-dependent threshold is reached where the processor performance for reducing the next increment of storage has too high a penalty. The amount of storage just short of this point is the new, reduced storage requirement.

Because of improvements to the scalability design in V4R4, the amount of virtual (disk) storage required to support a given thread has increased. *This does not always or necessarily typically imply an increase in main storage*. It is important to understand that AS/400 behaves differently than other machines. What the changes in V4R4 will do is make successful use of the GCHMAX parameter of the Java command more difficult. This is because GCHMAX in AS/400 Java does not necessarily control main storage consumption, nor did it even in V4R3. In V4R4, certain changes may magnify this existing effect. Since most people specifying -mx on a conventional Java or GCHMAX on AS/400 really wish to control main storage, the effect of the parameter on AS/400 Java may not match either intuition or V4R3 results.

Recall that we have recommended the usual value be GCHMAX(*NOMAX) even for V4R3.

Despite all this, some customers who used GCHMAX to control storage consumption in V4R3 may see little or no change, because the per-object cost for most Java objects is substantially reduced and will offset the new virtual storage "minimums."

Those who have a comparatively low load on a given Java virtual machine and who specify GCHMAX may see an apparent increase in requirements. Because storage consumption is so complex, even the improvements that most will witness may upset some other customers' existing Java command invocation strategies. While this will *appear* to cause problems such problems should ordinarily vanish with a bit of adjustment to the Java command.

Steps:

1. Run with your parameters on the Java command unchanged from V4R3. Most applications should see overall improvement in memory consumption.

2. If there is no improvement, or experimenting is desired, adjust GCHMAX. A simple fix would be to increase GCHMAX in the range of 5MB to 15MB. If this is not satisfactory, another approach would be to experiment with GCHMAX(*NOMAX) and use the size of the storage pool the Java application runs in to control main storage consumption. Because of the effect of Single Level Store, GCHMAX(*NOMAX) and a specified storage pool may be the best way of controlling main storage consumption. The JVM and Storage Management normally do a fine job of controlling storage. The overall effect of specifying

*NOMAX and a proper pool size can be more efficient in main storage usage than the -mx parameter of competitor's Javas, which directly controls both virtual and real storage.

3. If step 2 fails, and the application allows variability in the number of Java virtual machines and/or the number of threads, try running with fewer threads (and, possibly, more JVMs) and see if an improvement is achieved.



7.3 JDBC Improvements and Commercial Java Applications

One key and frequent distinguishing mark between strictly computational programs and commercial programs is use of a data base, usually intense use of a data base.

IBM has sponsored a benchmark called jBOB (Business Object Benchmark for Java) and is in the process of releasing it so that there can be independent validation of results. Since this benchmark is database oriented, it cannot be simply downloaded and run.

Independently validated results suggest that AS/400 can achieve a substantially better performance than NT (22%).

The above data, while not independently verified, is consistent with the independently validated results. This shows release to release improvements has trended roughly on the range of 35 percent compound improvement in the software path length.

Figure 7.2 Java Data Base Improvement

7.4 Unmatched Scalability



Figure 7.3 Java Scalability, 4-way AS/400, 4-way NT

AS/400 Java maintains its scalability advantages over the competition. Our published benchmark, jBOB, which should shortly become an industry sponsored and industry standard benchmark, shows AS/400's superior scalability.

Especially when the configurations are larger, an AS/400 will be an excellent choice because it will be able to scale farther and better that competitive offerings.

There are cases where the NT implementation will outperform AS/400 at smaller loads. Industry publications are sometimes based on these smaller, briefer, and easier-to-run evaluations. However, when these are scaled up, sometimes as little as by a factor of 10, AS/400's scalability advantages begin to show up.

The only platform we know of which continues to rival AS/400's scalability story is RS/6000 under AIX.

7.5 Comparison to Existing Languages

Despite improvements, Java's performance continues to lag more traditional languages in performance. Despite the dramatic gains in the last two AS/400 releases, Java remains behind other languages. That is because the early Java had so far to go. Still, the gap for Java versus other languages is now often under 2-to-1 in performance. History suggests that now that Java has breached this 2-to-1 barrier (and can even sometimes beat existing languages), that Java's other advantages will begin to carry the day. That is, once new programming technologies, especially those showing productivity advantages, are less than 2-to-1 behind their predecessor, the critical mass for long term acceptance has been reached. While the rate of improvement may now slow, Java's relative improvement compared to more mature languages should continue. It's general improvement pattern matches the historical tracking of other "revolutions" such as high level language versus assembler, structured over non structured, and C++ O-O over ordinary C. In these cases, the new technology initially had performance problems, but improved over its more mature rival to the point that any remaining deficit was ignored in favor of the other advantages.

The lag in performance, however, persists and for both data base intensive and non-data base intensive applications and servlets. Reasons for this continue to be Java's relative immaturity (Java acceptance has occurred at an unprecedented rate), and some common Java programming idioms which ease programming, but hurt performance. Garbage collection can often be costly as well, a cost other languages avoid at a significant software reliability cost. See "Tips and Techniques" below for ways to minimize these problems. Before going further, we should note that all competitors we are aware of continue to share this performance deficit.

Last year, we suggested industry projections showing Java already at or near parity with other languages remain largely hype. This is no longer as true as it was. There are beginning to be cases where Java outperforms other languages, but that holy grail remains elusive for large, significant applications.

Internal work has confirmed that typical computational Java programs will lag traditional languages. Typically, Java is slower by factors ranging from 1.75 to 3.

Java also lags in commercial performance. Java (via JDBC) uses dynamic SQL and extended dynamic SQL. SQL is very general and functionally rich, but has a certain overhead, especially for simpler functions. Traditional COBOL and RPG use simpler, highly tuned native data base access. Accordingly, data base access will not reach parity with traditional languages. For Java programmers willing to bypass a general Java solution and use native AS/400 data base, the AS/400 Toolbox for Java offers alternative methods for achieving better data base performance.

There will also be some cases where Java will beat RPG because it either avoids a sort or the need and expense of maintaining a keyed access to a given data base (Java can sometimes eliminate a sort by using an application-constructed object network).

That said, we expect most developers to use the standard JDBC access methods because it is part of the Java portability story.

Java's strengths in the V4R4 time frame will continue to be:

- Multi-platform vendors who wish to exploit the "write once, run anywhere" property and sell their applications on AS/400 as well as other Java machines
- Those needing simplified access to advanced functions such as networking, multithreading, client-server, and object-oriented programming,
- Those who wish to web-enable variations of existing applications (e.g. order entry becomes customer order via web),

- Those looking for simple, productive access to SQL function,
- Those needing to bridge from the ASCII world (either from IFS files or the network) to EBCDIC data bases.

Customers needing maximum application performance should look at traditional languages like COBOL or RPG. C++ is also available on AS/400, providing higher performance for Object-Oriented applications, but at a cost of greater complexity for programmers.

Some Java programmers may find that most of their performance problems occur on only a couple of highly used routines. For applications of this type, the Java Native Interface (JNI) allows the high use routines to be written in C or C++, but inter operate with the rest of the application in Java. If carefully limited to a few critical routines, better performance is gained without giving up Java's other advantages.

7.6 Java Performance -- Tips and Techniques

Introduction

Tips and techniques for Java fall into several basic categories:

- 1. AS/400 Specific. These should be checked out first to ensure you are getting all you should be from your AS/400 Java application.
- 2. Java Language Specific. Coding tips that will ordinarily improve any Java application, or especially improve it on AS/400.
- 3. Data Base Specific. Use of data base can invoke significant path length in OS/400. Invoking it efficiently can maximize the performance and value of a Java application.
- 4. Garbage Collection and Allocation Specific. Because Java programmers don't directly return their unused storage for reuse, the Java garbage collection facility must run on occasion to claim unused storage. Tuning the execution of garbage collection can be highly important to performance. This can be done by tuning garbage collection's performance or by avoiding the creation of new objects (see also language specific suggestions).

AS/400 Specific Java Tips and Techniques

- Load the latest CUM package and PTFs
 To be sure that you have the best performing code, be sure to load the latest CUM packages and PTFs
 for all products that you are using. Information on the AS/400 JVM can be found at the
 "http://www.softmall.ibm.com/as400/java/devkit.html" Developer Kit for Java Web Site. Information
 on the AS/400 Toolbox for Java can be found at the "http://www.as400.ibm.com/toolbox" Toolbox
 Web Site.
- Use CRTJVAPGM on .class files Java .class files should be converted into direct execution (machine instruction) Java program objects through the CRTJVAPGM command. Use this command before running any Java programs. The
Program object is permanent and will be reused once it is created. Since this program is implicitly created by Java/400, it is normally "hidden" and not visible. To see if it exists, use the DSPJVAPGM command. Optimization levels 10, 20, 30 and 40 are supported on the CRTJVAPGM command. After debug, in most cases, optimization level 40 should be used. Opt 40 ordinarily gives the best performance. Opt 20 might be a better choice for debugging.

Relative Performance (Optimization level):

Results of specifying a given optimization level will vary by application. For computation and call intensive applications the relative gains can be dramatic. Here are the relative performance gains for a well-known computational problem, "The Towers of Hanoi."

	Relative time (bigger is slower)
Optimization level 40	1.00
Optimization level 30	1.73
Optimization level 20	2.02
Optimization level 10	2.65
Interpretive	12.95

Comparisons based on V4R2, JDK 1.1.4. Similar magnitudes would be observed on V4R3 and V4R4.

• Use CRTJVAPGM on .zip and .jar files

CRTJVAPGM should be used on .zip and .jar files. The JAVA/RUNJVA command will implicitly create program objects for each class file in a .zip or .jar. While the program created implicitly is now retained (in V4R3 and V4R2 it was discarded), the optimization level for the implicit creation case is only optimization level 10. Because of certain optimizations in V4R4, the difference between optimization level 10 and optimization level 40 can be even greater for a .jar or a .zip than in the .class case shown above. The implicit creation of program objects will also significantly impact JVM startup or class loading. With very large .jar or .zip files, this can look like a hang. To determine if your .zip/.jar file has a permanent program object use the DSPJVAPGM command. *Note:* The retaining of the program created for the .jar or .zip file can increase auxiliary storage requirements over V4R3. Normally, this should be tolerable, but it may occasionally be noticed.

- *Package your Java application as a .jar or .zip file.* Packaging multiple classes in one .zip or .jar file should improve class loading time and also code optimization in V4R4. Within a .zip or .class file, AS/400 Java will attempt to in-line code from other members of the .zip or .jar file.
- Use the special property os400.defineClass.optLevel for dynamically loaded .classes Java's definition will occasionally cause the results of CRTJVAPGM to be ignored. This is especially true if your Java program loads a class "by hand" (Class.forname(), ClassLoader.loadClass()). In these cases, Java/400 cannot know the name of the file from which the class came, (strictly speaking, there may not be a file) so it must decide between interpretation and class loading using only the byte array provided by the defined interfaces. The os400.defineClass.optLevel property, which can be passed as a property through the Java command, will tell Java/400 whether to interpret or compile the program. Remember to pass the name and optimization level properly:

JAVA CLASS(your.main.class) PROP((os400.defineClass.optLevel 40))

It is often beneficial to set this to optimization level 20. Note: Servlets will need to have this property passed via configuration options in your web server, because the web server invokes the Java command. Consult the appropriate product documentation.

• *Be aware of some automatic recreation of "hidden" programs in V4R4.*

Java/400, to improve performance, is changing the internal format of the hidden *PGM object created by CRTJVAPGM. All existing V4R3 *PGM objects will be recreated on their first use at the same optimization level as in V4R3 and become V4R4 objects unless someone uses CRTJVAPGM on the .class, .jar, or .zip file before the first use.

If no action is taken, no harm is done; the recreation of the hidden program will commence. However, this change means the first use of a Java class in V4R4 that was unchanged from the V4R3 migration may appear to run more slowly. If you do the CRTJVAPGM yourself at a relatively benign time, this overhead should not affect production use of your machine even this one time. Doing the CRTJVAPGM by hand before use will be particularly beneficial for .zip and .jar files. It also means that if your program runs slowly in V4R4, try it again and see if the slowdown goes away. If it does, some class probably underwent compilation for migration. *Note:* This is strictly a performance issue. You *do not need* to recompile your .java source or make any other changes to your program because of this activity. The classes shipped with OS/400 JV1 are already at the V4R4 level.

Java Language Performance Tips

• Minimize synchronized methods

Synchronized method calls take at least 10 times more processing than a non-synchronized method call. Synchronized methods are only necessary if you have multiple threads sharing and modifying the same object. If the object **never** changes after it is created ("constructed" is the Java term for "created"), you don't need to synchronize any of its methods, even for multithreading.

Note: Dealing with synchronized methods mean understanding some important Java programming concepts.

- Some Java objects, notably String, do not permit data in the object to be modified after the object is constructed. For such objects, synchronized methods are never needed.
- Other objects, such as StringBuffer, allow the object to be modified after construction. All of its methods are synchronized.
- Many objects fit these two models. If a StringBuffer type object will be used by even one multithreaded application, all methods must be synchronized except its constructors. If you never use multithreading, then a StringBuffer type object requires no synchronization.
- But, consider object reuse when you decide. If some later application uses your object, and that new application is multithreaded, synchronization will be needed. This is why common Java objects like StringBuffer have synchronized methods.
- Use the -O javac option

The javac compiler may in-line methods if the -O option is selected. In order for a method to be in-lined it must be a final method (static, final, or private). Final methods can not be subclassed, which is the

key to allowing in-lining. AS/400 Optimization 40 in V4R4 now includes some selective in-lining without specifying -O.

• Minimize object creation

Object creation can occur implicitly within the Java APIs that you use as well as within your program. Object creation and the resulting garbage collection can typically take 15% to 30% of a server transaction workload. To minimize this cost you can reuse an object's space implementing a "reset" method that reinitializes the local variables in the object. The code fragment

```
if objx = null then
   objx = new x(some,creation,parameters);
else
   objx.reset(some,recreation,parameters);
```

can provide significant performance improvements.

Common causes of object creation that may not be obvious:

- The I/O function readLine() creates a new String.
- Invoking the substring() function of a String creates a new String.
- The JDBC Result Set function getString() creates a String.
- The StringTokenizer returns a String from many functions.
- Passing a scalar int or long as an object will create an Integer or Long object.

• Minimize the use of String objects

String objects in Java are immutable. This means that you can not change (append, etc.) to a string object without creating a new object. Object creation is expensive and can occur multiple times for each String object you are using. To minimize the use of String objects you should use either StringBuffer or char[]. StringBuffer may also be a problem since the StringBuffer classes use synchronized method calls. An array of characters (char[]) can be used to simulate fixed length strings. This is recommended for applications which make heavy use of string data.

Relative Performance:

The following table shows the relative performance difference when using String, StringBuffer, or char[]. The test case concatenates two strings. For the char[] case. the concatenation reduces to simple array assignment, thus avoiding the creation of objects and the synchronization overhead associated with StringBuffer. In the following table the String "World Wide" was concatenated to the string "Wait". For the char[] case there were simply four char[] assignments for the characters 'W' 'A' 'I' 'T'.

	Relative time (bigger is slower)
char[]	1
StringBuffer	35
String	363

Comparisons based on Optimization level 40, V4R2. JDK 1.1.4. V4R4 results would be similar. Leverage variable scoping Java supports multiple techniques for accessing variables. One typical technique is to write an "accessor" method. Local variables are the fastest.

• Leverage variable scoping

Java supports multiple techniques for accessing variables. One typical technique is to write an "accessor" method. Local variables are the fastest.

Relative performance:

Here are five comparisons on variable access time and their relative performance. A local variable is the fastest and is given a relative performance of 1.

	Relative time (bigger is slower)
Local variable	1.0
Instance variable:	2.0
Accessor method in-lined:	2.0
Accessor method:	5.9
Synchronized accessor method:	99.6

Comparisons based on Optimization level 40, V4R2. JDK 1.1.4. V4R4 relative differences would be comparable.

Note: This is a performance-oriented suggestion. Making instance variables public reduces the benefit of Object Orientation. Having a local copy in the method of an instance variable can improve performance, but may also add coding complexity (especially in cases where individual blocks use the synchronized keyword). Avoiding the "synchronized" label on a method just for performance may lead to difficult bugs in multithreaded applications.

• *Minimize use of exceptions (try catch blocks)*

Exception handling requires a certain amount of "setup" information. A "try" statement in Java results in some setup so that if an exception occurs, the "catch" routine will get control. This adds path length.

Sometimes, exceptions are required by a given method on a given object. These must obviously be coded. But, try and avoid having too many of these.

• Do not start up too many JVMs

The JAVA/RUNJVA commands create a new batch immediate Job to run the JVM. Limit this operation to relatively long running Java programs. If you need to invoke Java frequently from non-Java programs consider passing messages through an AS/400 Data Queue. The ToolBox Data Queue classes may be used to implement "hot" JVM's.

• Use static final when creating constants When data is invariant, declare it as static final. For example here are two array initializations:

```
class test1 {
      int myarray[] =
                 \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, \}
                          2,3,4,5,6,7,8,9,10,11,
                          3,4,5,6,7,8,9,10,11,12,
                          4,5,6,7,8,9,10,11,12,13,
                          5,6,7,8,9,10,11,12,13,14 };
      {
class test2 {
      static final int myarray2[] =
                  \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, \}
                                          2,3,4,5,6,7,8,9,10,11,
                                          3,4,5,6,7,8,9,10,11,12,
                                          4,5,6,7,8,9,10,11,12,13,
      }
                                          5,6,7,8,9,10,11,12,13,14 };
```

Relative Performance:

When 10,000 objects of type test1 and test2 were created the relative time for test1 was about 2.3x longer than test2. Since the array myarray2 in class test2 is defined as static final, there is only ONE myarray2 array for all 10,000 instances of test2. In the case of the test1 class, there is an array myarray for each test1 instance.

Comparisons based on Optimization level 40, v4r2. JDK 1.1.4.

AS/400 Database Access Tips

```
• Use the native JDBC driver
```

There are two AS/400 JDBC drivers that may be used to access local data. Programmers coding connect statements should know that the Toolbox driver is located at Java URL *"jdbc:as400:system-name"* where system-name is the AS/400 TCP/IP system name. The native JDBC driver is located at Java URL *"jdbc:db2:system-name"* where the system-name is the Data Base name. The native AS/400 JDBC driver uses an internal shared memory condition variable to communicate with the SQL/CLI Server Job. The ToolBox JDBC driver assumes that the data is remote and uses a socket connection into the client access ODBC driver. The native JDBC driver is faster when you are accessing local data.

• Use Prepared Statements

The JDBC prepareStatement should be used for repeatable executeQuery or executeUpdate methods. If prepareStatement is not used, the execute statement will cause an implicit prepareStatement every time the execute statement is run.

Note: Avoid placing the prepareStatement inside of loops (e.g. just before the execute). In some non AS/400 environments, this just-before-the-loop coding practice is common for non Java languages. Programmers may carry it over to Java. However, in many cases, the prepareStatement doesn't change and the Java code will run faster on all platforms if it is executed only one time, instead of once per loop. It will show a greater improvement in AS/400.

• Store character data in DB2 as Unicode

The AS/400 JVM stores string data internally as 2 byte Unicode. If you are reading or writing large amounts of string data and the data is stored in EBCDIC, the data will be converted on every database access. You can avoid this conversion by storing the data in DB2 as 2 byte Unicode. Use the SQL

graphic type with CCSID 13488 or the DDS graphic type with CCSID 13488.

Note: Be careful with this suggestion. 1) If characters are the main portion of the record data, the record could double in size. If this is a large and important data base, this will increase hard disk expense, perhaps by a large amount. 2) If the data base is accessed by non Java code (e.g. legacy RPG applications) Unicode may create complications for the old code.

• Store numeric data in DB2 as double

Decimal data cannot be represented in Java as a primitive type. Decimal data is converted to the class Java.lang.BigDecimal on every database read or write. This conversion can be avoided by storing numeric data in the database as float or double. Alternatively, you may considering converting the BigDecimal type to double for heavy calculations while leaving the Database in decimal form. Be aware that rounding differences may be introduced with the use of double. Float can also be used, but most C programmers use double for this purpose to minimize rounding difficulties. In rare cases (e.g. in the banking industry) decimal math can be a requirement. Use BigDecimal for these.

• Use ToolBox record I/O

The AS/400 ToolBox for Java provides native record level access classes. These classes are specific to the AS/400 platform. They provide a significant performance gain over the use of JDBC access.

• Check ToolBox for existence of a Java program object

The jt400.jar file contains the AS/400 ToolBox for Java product. After installation, this .jar file does not have a Java program object. Use the CRTJVAPGM at optimization level 40 to create the program object. Use the CRTJVAPGM command during low system activity as it will take some time. Use the DSPJVAPGM command to see if the program object already exists.

Allocation and Garbage Collection

• Leave GCHMAX as default

The GCHMAX parameter on the JAVA/RUNJVA command specifies the maximum amount of storage that you want to allocate to the garbage collection heap. In general the default value (Set to the largest possible value) should be used. The system does not allocate the storage until it is needed. A large value does not impact performance. If this maximum is reached, the JVM will stop all threads and attempt to synchronously collect objects. If GCHMAX is too small, a java.lang.OutOfMemory error will occur. Some V4R4 improvements may rarely cause difficulties if GCHMAX was small in V4R3, but successful there.

• Adjust GCHINL as necessary

The GCHINL parameter on the JAVA/RUNJVA command specifies the amount of initial storage that you want to allocate to the garbage collection heap. This parameter indirectly affects the frequency of the asynchronous garbage collection processing. When the total allocation for new objects reaches this value, asynchronous collection is started. A larger value for this parameter will cause the garbage collector to run less frequently, but will also allocate a larger heap. The best value for this parameter will depend on the number, size, and lifetime of objects in your application as well as the amount of memory available to your application. Use of OPTION(*VERBOSEGC) can give you details on the frequency of garbage collection, and also object allocation information.

- *Ignore GCHPTY* This parameter is not used. It has no effect on performance.
- *Ignore GCHFRQ* This parameter is not used. It has no effect on performance.
- Monitor GC Heap faulting

Java objects are maintained in the JVM heap. Excessive page faults may occur if the memory pool for your JVM is too small. These faults will be reported as non-database page faults on the WRKSYSSTS command display. Typically, the storage pool for your JVM is *BASE. Fault rates between 20 and 30 per second are acceptable. Higher rates should be reduced by increasing memory pool size. In some cases, reducing this value below 20 or 30 per second may improve performance as well. If you have the storage available, reducing the rate below 20 to 30 per second may be a benefit. Lowering the GCHINL parameter might also reduce paging rates by reducing the AS/400 JVM heap size.

• *Minimize Object Creation* See previous suggestion about minimizing object creation.

7.7 Capacity Planning and Model Recommendations

Java requires more resources than previous languages. Accordingly, when estimating capacity, a more robust machine should ordinarily be specified.

Java's added resources have been diminishing over time (see the previous sections of this chapter).

Still, all in all, it costs more resource to deploy Java than a traditional RPG application today.

Things to consider when estimating a machine with Java content:

V First, remember to account for the amount of traditional processing (RPG, COBOL, etc.)

going on. To the extent traditional work is going on, the machine should be sized according to existing capacity planning guidelines. This also applies to customers using newer functions like SAP and Domino whose characteristics are well-known. If Java is present in some applications, but most of the machine cycles are spent on traditional work, or (say) SAP content whose machine requirements are known, Java might fit in with little or no adjustment to capacity planning. The main caveat then becomes Java's growth rate versus traditional applications. If Java is the core of someone's e-business and e-business grows rapidly, so will their Java content.

▼ Second, be careful to ascertain how much Java is going into the AS/400 itself. If AS/400 is being accessed by client Java code, but the code in the AS/400 itself remains in COBOL, RPG, or C/C++, there's no point in padding capacity for Java -- it isn't being used. The important item in that case becomes ensuring the customer's PCs have enough performance to run Java on the client, and enough traditional horsepower to service their requests. Likewise, if the AS/400 is just being used as a Web Server (e.g. Domino GO), there's no need to change capacity planning for Java content for that reason alone. Until Java is used for servlets, an AS/400 running a Web Server will not itself be running Java function. The capacity needed to run the Web Server is instead the main issue.

- V Third, even when Java servlets and Java applications are being used, account carefully for added system services. Web serving, communications, and data base costs can often swamp Java's contribution to an end-to-end application. Because it uses JDBC and dynamic SQL, Java can increase the data base costs compared to a traditional application doing similar things.
- V Fourth, recognize that AS/400 has been optimized around scaleable, OLTP type applications which use lots of system services such as data base. Java, by contrast, will tend to put more of its execution in the application itself. In the short run, simpler servlets may complicate this story, but over time, Java content will grow as a percentage of the processor compared to traditional. The reason relates to Java's portability story. Java will tend to invoke Java-based function where RPG would invoke the operating system. This property will tend to increase processor requirements overall compared to what we're familiar with. Other features of Java will tend to require more main storage than traditional languages.

V Fifth, because of the increased processor needs, be wary about using the smallest AS/400 models. This is particularly true of test and development machines. Because of OLTP price/performance tradeoffs, smaller or older machines may be disproportionately disappointing to customers when used for Java, even for testing. In general, make sure the processor performance of the test machine and development machines is in line with that of the production machine. This would mean deploying a machine with a higher uniprocessor CPW rating than would ordinarily be the case.

V Sixth, beware of misleading benchmarks. Many individuals will be willing and able to write their own benchmarks for Java. They'll also be able to download some "Java benchmarks" over the Internet. *Most of these will be poor predictors of server performance*. This includes VolanoMark, which requires careful tuning and primarily measures Communications Performance. Because of this, and Java/400 tradeoffs for better server performance, many of these sorts of benchmarks will also tend to make AS/400 look worse than the actual deployment of their application would be. Those running a Java evaluation should make sure that any benchmark: a) is some kind of prototype of a true 'server' application, b) runs long enough (at least 15 minutes) to represent a fair, steady-state comparison, c) has scalability characteristics (multiple threads, multiple Java jobs, etc.). AS/400 Java is not optimized for simple, single-threaded benchmarks. Nor should it be: AS/400 customers will tend to deploy multiple servers and threads in a typical Java use (e.g. web serving via servlets). Another thing to watch for: Using an inadequate test machine for bench marking and then fear Java isn't acceptable on their bigger, faster production one.

V Seventh, recognize that Java won't deploy in a traditional manner. 5250 operations to and from Java will not be a frequent attribute of Java on the AS/400. Accordingly, the higher the Java content, as a percentage of total operations, the more a server model should be considered. With the new V4R4 model line, a lower interactive rating should be considered as Java content increases.

V Eighth, keep in mind that not everything changes for Java.

1. Whether SMP (Symmetric Multiprocessing) makes sense will not typically change for Java. Java probably will run better with a machine using the fewer CPUs for the same CPW rating, but this is very often true of traditional applications as well.

2. The size of the data base for Java should be about the same. Since data base often swamps other uses of DASD, that means that Java should seldom require more disk space than traditional languages. If Java's performance is low enough compared to RPG (and, with JDBC's dynamic SQL, it can be) DASD cost could even be less.

The March Of Time

It is difficult to write a single set of guidelines that can successfully set expectations for everyone. Customers who have bought machines fairly recently, and have a while to go before the budget allows an upgrade, may be less interested in comparing performance with current products than those whose identical machine is considered near the end of its useful life.

Meanwhile, new AS/400 and competitive machines pour into the marketplace. The improvement in price/performance, even in a single year, can be quite steep. It is sometimes a shock to compare what one owns with what is new.

This set of recommendations is largely written from the point of view of someone wishing to purchase a machine or make decisions about upgrading an older machine. Such persons will have a lower tolerance for performance problems than someone who happily use their existing machine, with no thought of replacement.

Better Machines for Java

The following machines are more likely to give satisfaction, relative to their CPW rating, when used for Java. For the new V4R4 machines, there is no distinction made here between (say) a 2064 with interactive feature code 1502 and feature code 1505. Generally, that is an issue that is decided based on the ratio of client/server, batch, and traditional (RPG and COBOL) interactive work. Java servlets, client/server applications or batch applications all tend to be charged as "batch" work. Thus, the interactive feature code is something to decide after one has sized the machine generally for both its Java and traditional content.

AS/400 Model	AS/400 Feature
740	2070, 2069
730	2068, 2067, 2066, 2065
720	2064, 2063
170	2388, 2386, 2385
650	2243, 2240
640	2239, 2238
530	2162, 2153, 2152
S40	2261, 2256
S30	2260, 2259, 2258
S20	2166
538	2157, 2156, 2155

Machines to Be Used With Caution with Java

The following machines can be used with care. Existing owners of these machines may be quite content with the performance. For new deployments, they are especially suitable when Java is present, but Java's

actual contribution to the total execution is low (say, 50% or less). Simple servlets, for instance, with minimal data base function may do very well on these machines. The higher the traditional content, the more acceptable these machines will be. These machines will also be more acceptable in cases where an existing V4R3 Java application on the AS/400 is already deployed and, therefore, more precise estimates and specific requirements are available:

AS/400 Model	AS/400 Feature
720	2062, 2061
170	2292
620	2182
S20	2165

Machines to Be Used With Extreme Caution for Java

Some customers will find that their Java application will work well with any AS/400 whatever. That said, in a guideline like this, one must turn toward expected averages and not specific cases. Because of Java's general tendency to consume more resources than traditional applications, and because the margin of error is correspondingly less with these machines, care must be taken in their use for Java. Many Java applications simply won't perform well on the machines on this "Extreme Caution" list. On the other hand, some will find, despite the caution expressed here, that some Java applications can be deployed on these machines. This is especially true if Java's overall contribution can be limited compared to other work (e.g. 25% or less of total workload).

Many Java applications will not work well for the machines on this list unless:

1. The application is well-understood in terms of its processor consumption requirements so that these machines are usable. For instance, previously running the application on a different V4R3 AS/400 machine may permit reasonable extrapolations for a proposed new machine. In some cases, disk and Communications I/O performance might need to be well-understood also.

2. The customer is able to estimate capacity requirements when initially deploying the application and over a reasonable period of time. If the application is well known to support (say) 10 users for a particular model and feature code, there's no point in using that feature code if 20 users are expected reasonably soon. This idea is obvious, but since Java typically consumes more resources, any excess capacity will disappear more quickly as growth occurs than one might otherwise expect.

3. In some cases, the customer deploying the application must be able to upgrade to at least V4R3 and probably V4R4. Since Java performance has improved rapidly, OS upgrades are an obvious way to improve the apparent performance of all AS/400 machines and these in particular.

The 170 models 2291, 2290 and 2289 are aimed at environments with traditional workloads. Like any machine in this list, they can be used for Java, in a functional sense. However, they are not designed with Java's higher consumption in mind and so may disappoint, especially in the context of a newly purchased machine.

The 170 features 2183, 2176, 2164, 2160, and 2159 are all older by up to two years than the newer 170s (such as the 2385) and have an older, slower processor. Such machines might still give their current

owners satisfaction, but cannot compare in Java performance to the newer AS/400s or current competition.

The 150, 600, 510, 500, 50S, 400, 40S, and S10 lines represent older AS/400 processors. Java was built with a newer and faster world in mind than these boxes provide. Moreover, the march of time has reduced people's tolerance for the performance of these older machines in a Java context.

AS/400 Model	AS/400 Feature
170	2291, 2290, 2289, 2183, 2176, 2164, 2160, 2159
640	2237
620	2181, 2180, 2179, 2175
600	2136, 2135, 2134, 2129
530	2151, 2150
510	2144, 2143
500	2142, 2141, 2140
400	2133, 2132, 2131, 2130
S30	2257
S20	2163, 2161
S10	2119, 2118
538	2154
50S	2122, 2121, 2120
40S	2112, 2111, 2110, 2109
150	2270, 2269

Relative to their CPW rating, and sometimes in an absolute sense, the following machines are more likely to cause customer dissatisfaction when used for Java:

Chapter 8. IBM Network Station Performance

Performance information for Releases 1 to 3 for IBM Network Stations attached to V3R2, V3R7, V4R1, V4R2, V4R3 and V4R4 is included below. The following IBM Network Station functions are included:

- Time to initialize the IBM Network Station (prior to login) for ethernet, token ring and twinax
- Time to load the applications (5250 emulation and browser, etc.)
- 5250 application performance
- Browser performance
- Java Virtual Machine (VM) applet/application performance
- Times for the IBM Network Station series 100, 300 and 1000

The computer industry has a generic name for the IBM Network Station - the thin client. Since clients attach to servers, it might seem that an AS/400 SERVER model attached to a Network Station (a thin CLIENT) would always be the best fit (that is, CLIENT to SERVER). Be cautious when the Network Station is attached to an AS/400 SERVER model. When using 5250 applications, the Network Station looks like a Non-Programmable Terminal (NPT) (like an interactive job) to the SERVER and will be subject to the AS/400 SERVER interactive rules, so it might not always be a good fit. The traditional AS/400 SYSTEM models are always a good fit with the Network Station.

In the following sections, references to Release x refer to the Network Station. References to VxRx refer to the AS/400 releases. In addition, in the following sections, performance data for one Network Station is real data; performance data for 10, 50 and 100 Network Stations is simulated. All twinax data is real.

8.1 IBM Network Station Network Data

The following table shows the amount of data that flows from the AS/400 to each IBM Network Station for initialization and each application load:

Table 8.1. Elements Loaded to a Network Station (MB)					
Release	Rel 1-2.5	Rel 1-2.5	Rel 3	Rel 3 DBCS*	
Series	100/300	1000	All	All	
Kernel +	4.0	4.8	3.0	3.9	
Configuration +					
Other					
5250 Emulation	0.9	0.9	1.6	3.8	
3270 Emulation	0.3	0.3	0.9	3.2	
IBM Network	2.2	2.2	NA	NA	
Station Browser					
Navio NC Navigator	3.7	3.7	5.0	10.0	
Java Virtual Machine	1.5 - 5.0	1.5 - 5.0	1.5 - 5.0	1.5 - 5.0	
Note:					
• The amount of data downloaded will vary, depending on the configuration selected					

• *DBCS support includes Japanese, Korean, simplified Chinese, and Traditional Chinese

The kernel/configuration data is downloaded when the Network Station is powered-on. Unless configured otherwise, all the other options are downloaded when they are selected.

Note that when an application (e.g. 5250 emulation) is closed or the user logs-out, that application will again be downloaded when it is next selected - it is not kept in memory across log-outs. The kernel/configuration data is kept in the Network Station across log-outs.

The Java Virtual Machine download time varies depending on the application. Only the required classes are downloaded.

In Release 3, some of the information that is sent to the Network Station is compressed. Once received, the Network Station decompresses it. This compression means fewer bits are shipped from the AS/400 to the Network Station, resulting in better LAN utilization. More data/function is shipped to the Network Station in Release 3 than in previous releases. The compression results in boot performance that is about equal to previous releases.

Release 3 contains an option, TFTP subnet broadcast, that can significantly decrease the amount of data transmitted during the boot process, as well as saving significant CPU cycles in the AS/400. This option is described further in the sections below.

8.2 IBM Network Station Initialization

Initialization, at this time, is not trivial and could be a performance concern for some customers. The time to initialize the Network Station, particularly when many stations are initialized simultaneously can be prohibitive. In addition, initialization can consume a lot of AS/400 CPU, so that other jobs on the AS/400 might be starved.

If possible, it is best to leave the Network Station powered-on after initialization and/or to stagger initialization. The IBM Network Station consumes very little power. If initialization times are a problem and power outages are a concern, battery backups for each IBM Network Station, should be considered or possibly server systems dedicated to initialization.

Different Initialization Mechanisms - the Gory Details

Initialization is performed using TCP/IP Trivial File Transfer Protocol (TFTP) and/or AS/400 Remote File System (RFS). Both these access methods read files from the AS/400 to the Network Station. For reliability and performance, both mechanisms subdivide files into blocks for sending, and then recombine them in the Network Station. The TFTP block size can be configured 512 thru 8192 bytes. The RFS block size is fixed at 8192. For Releases 1-2.5, TFTP and/or RFS is used during initialization depending on the configured initialization options. For Release 3, the kernel is loaded using TFTP and then RFS is automatically tried - no matter what the configuration - if RFS is unsuccessful, the configured options are tried.

There are three possible ways to initialize the Network Station:

• NVRAM - the AS/400 and Network Station IP addresses and other information are configured in each Network Station. The Network Station sends a TFTP request to the configured server to begin initialization.

- BOOTP the Network Station broadcasts to find a responding AS/400 server. The AS/400 server is previously configured with each Network Station's IP address and other information. Once the server receives a broadcast from a Network Station, it sends the configured data to the Network Station and then begins the initialization.
- DHCP the same as BOOTP except the AS/400 server contains a pool of Network Station IP addresses.

BOOTP or DHCP is the preferred method, for Releases 1-2.5. All methods are OK in Release 3.

For Releases 1-2.5, NVRAM uses TFTP to load the kernel/configuration files and, after login, uses RFS. For Release 3, NVRAM uses TFTP to load the kernel and RFS for all subsequent files.

BOOTP and DHCP use TFTP to load the kernel and then use RFS to load all subsequent files.

For Releases 1-2.5, the Network Station tries 10 times with a 5 second timeout to locate and read the kernel using TFTP. After 10 attempts, an error message is issued. For Release 3, the number of reties can be configured - an infinite retry is preferred.

For Releases 1-2.5, if NVRAM is selected, the Network Station reads the configuration files using TFTP. The Network Station will try 10 times with a 3 second timeout to read each file. If unsuccessful, it will skip that file and then try to read the next file - which eventually results in an unsuccessful initialization. (RFS will not skip files.) From a reliability perspective, this makes NVRAM, for Releases 1-2.5, the least preferred booting mechanism.

Release 3 contains a new option - subnet broadcast. Subnet broadcast is supported on ethernet, token ring and twinax. When this option is selected, TFTP data (the kernel - about 2MB), is broadcast from the AS/400 server to any requesting Network Station. That is, the kernel is sent one time so that each Network Station receives it. When subnet broadcast is off, the kernel is sent individually to each Network Station, which means a lot more data on the LAN/twinax. The broadcast is only to a subnet (e.g. any Network Station on a single ring, such as 9.5.112.x). When Network Stations from different subnets request the kernel, the AS/400 provides a broadcast to each subnet. The data below shows that subnet broadcast uses less AS/400 CPU. Subnet broadcast is the preferred boot option (twinax has some special considerations mentioned below). There is a caution - some routers do not support broadcast and broadcast can cause other problems, if not configured properly.

Subnet broadcast is supported on twinax. Unlike ethernet and token ring, the twinax protocol does not support broadcast. What this means for twinax, is that, when subnet broadcast is selected, each frame is sent individually to each device. When all devices are expecting the broadcast, this option works well (meaning less AS/400 CPU is used). When all devices are not expecting the broadcast, this option is not great (more data on the twinax cable). The data below illustrates this. In general, customers should not use subnet broadcast for twinax.

Some customers who have Series 1000s have experienced performance problems. The Series 1000 supports both full duplex and half duplex. In general, the performance problem is caused by a configuration error. The Series 1000 tries to operate in full duplex mode, but a router or something else in the network supports only half duplex. The Series 1000 almost continuously runs into collisions on the Ethernet which will result in extremely slow performance.

Some customers who have token ring network switches that pass 4K frames have experienced difficulties. The customers had set their LAN frame size/MTU to a value greater than 8K. In general, these customers used NVRAM - with default 1024 TFTP block size. Initialization works fine until login - where RFS takes over and uses 8K frames. The 8K frames do not pass through a 4K switch. Some solutions to this problem might be: configure the switch to allow 8K frames, replace the switch with a router, or configure the AS/400 LAN frame size/MTU to 4K (twinax is fixed at 4K).

If the network has no Domain Name Server (DNS), performance can be very slow. The initialization logic expects a DNS. If none exists, initialization waits for DNS searches to timeout (30 second) before proceeding. AS/400 V4R2 contains DNS support. If a customer does not wish to use a DNS, for Release 3, good performance is still possible by doing the following:

- CFGTCP, Change TCP/IP domain information (option 12), set search priority to *LOCAL
- CFGTCP, Work with TCP/IP host table entries (option 10), add the IP address and host name for the AS/400 and each Network Station
- IBM Network Station Manager, select Hardware, select Workstations, under Domain Name Server, set Update Network Station Manager DNS file

The initialization options described in this white paper will fit most customer environments. There are other variations that can occur. For example, if the customer chooses BOOTP and successfully loads the kernel, but for some reason RFS isn't working properly, initialization will timeout on RFS and switch back to TFTP. Variations such as these are not described in this document. The BOOTP boot sequence is described in greater detail in the following section.

BOOTP Initialization

There are four steps in the BOOTP initialization process. To get a total initialization time, times from each of the following four steps must be added together:

1. Hardware Test

The hardware test is just that - a memory test and other hardware tests to ensure that the hardware is operational. For the most part, the length of this test is determined by the amount of memory in the IBM Network Station.

Table 8.2. Time (Seconds) to Perform Hardware Test					
Memory (MB)	Series 100	Series 300	Series 1000		
8	15	14			
16	18	18			
32	24	22	10		
48	30	26			
64	36	31	13		

2. Kernel/Configuration Initialization

In this step, the Network Station locates the AS/400 server, reads the kernel and configuration files, and then displays the login window.

The Network Station broadcasts a BOOTP request to locate the AS/400 server. Then the kernel (about 2MB) is read using the TFTP function of TCP/IP. And then configuration files are read using the Remote File System (RFS).

The time to load the kernel using TFTP is heavily dependent on:

- TFTP block size
- TCP/IP maximum transmission unit (MTU) size
- LAN line description frame size (fixed for twinax)
- TFTP subnet broadcast option
- number of TFTP jobs

The Network Station negotiates the TFTP block size with the AS/400. It can range from 512 to 8192 bytes. The Network Station default is 8192. In general, the Network Station uses the TFTP block size, MTU and frame size defined by the AS/400.

The AS/400 default TFTP block size is 1024. As will be seen in the following tables, best performance is obtained with a large TFTP block size (e.g. 8192). If the MTU or frame size are less than 8192 (e.g. Ethernet has a maximum frame size of 1492) performance can still be enhanced by configuring the block size greater than the MTU/frame size. If the TFTP block size is greater than the MTU/frame size. The Network Station TCP/IP fragments (subdivides) the TFTP blocks to fit into the MTU/frame size. The Network Station TCP/IP recombines the MTU/frames into TFTP blocks. This fragmentation provides better performance than setting the TFTP block size equal to the MTU/frame size. Users should be aware that some routers, switches and/or gateways do not support this fragmentation capability. Twinax MTU/frame size are fixed, so fragmentation does not apply to twinax attached Network Stations.

The number of TFTP jobs on the AS/400 is also a performance factor - the optimal number for a system with a single LAN IOP is about 6, the default. The TFTP jobs are a pool of AS/400 jobs that download the kernel to Network Stations. They are first come, first serve. If there are more Network Station requests than jobs, the excess are ignored (i.e. not queued). If a request is not satisfied, the Network Station, every 5 seconds, will repeat its request. In general, there should be 6 TFTP jobs for each LAN IOP that has attached Network Stations.

The following tables and figures show how the TFTP block size affects the kernel/configuration initialization time, for a few AS/400 system sizes. The tables also show what happens when 1, 10, 50, and 100 Network Stations simultaneously (e.g. after a power outage) request TFTP initialization. The times represent the number of seconds when the last Network Station completes its TFTP and RFS download.

The data in the following tables was obtained in a dedicated environment. That is, only BOOTP, TFTP and RFS were running on the AS/400 and there was no other load on the LAN. In each test case, the base pool (memory) was cleared before beginning the test.

Results listed here do not represent any particular customer environment. Actual performance may vary significantly from what is provided here.

Table 8.3. Kernel/Configuration - AS/400 F97 and Network Station 300							
	Kernel/Configuration Initialization Time in Seconds						
(Average CPU Utilization in %)							
	AS/400 Model F97 (V3R2)						
	IBM	I Network Station S	eries 300 (Releases 1	-2.5)			
		16Mb To	oken-Ring				
		8KB MTU/Frame	Sizes, 6 TFTP Jobs				
		Vary TFTI	P Block Size				
# NS	512	1024	2048	4096	8192		
1	109 (5.0)	46 (5.5)	34 (4.2)	29 (2.9)	26 (2.6)		
10	10 225 (27.0) 105 (31.0) 77 (22.6) 63 (17.1) 57 (12.2)						
50	50 992 (32.8) 470 (41.7) 327 (30.8) 257 (24.0) 209 (20.1)						
100	1885 (35.2)	890 (46.3)	624 (33.6)	503 (25.5)	395 (22.3)		
Nata							

Note:

• Results may differ significantly from those listed here

Table 8.4. Kernel/Configuration - AS/400 150-2270 and Network Station 300							
	Kernel/Configuration Initialization Time in Seconds						
		(Average CPU)	U tilization in %)				
		AS/400 Model 1	50-2270 (V3R7)				
	IBM	I Network Station S	eries 300 (Releases 1	-2.5)			
		16Mb Token	Ring MFIOP	,			
		8KB MTU/Frame	Sizes, 6 TFTP Jobs				
		Vary TFTI	P Block Size				
# NS	512	1024	2048	4096	8192		
1	85 (23.3)	35 (28.6)	31 (22.0)	27 (16.3)	26 (14.8)		
10	10 229 (87.8) 126 (82.2) 83 (72.9) 63 (63.4) 55 (53.6)						
50 1065 (94.2) 565 (95.0) 347 (92.0) 234 (87.6) 193 (77.6)							
100	2075 (97.5)	1119 (97.0)	682 (94.5)	448 (92.5)	352 (88.1)		
Note:							

• Results may differ significantly from those listed here

Table 8.5. Kernel/Configuration - AS/400 510-2144 and Network Station 300							
	Kerne	el/Configuration Init	ialization Time in S	econds			
		(Average CPU)	Utilization in %)				
		AS/400 Model 5	510-2144 (V3R7)				
	IBM	I Network Station S	eries 300 (Releases 1	-2.5)			
		2619 16Mb Toke	en-Ring LAN IOP				
		8KB MTU/Frame	Sizes, 6 TFTP Jobs				
	-	Vary TFTI	P Block Size				
# NS	512	1024	2048	4096	8192		
1	71 (9.8)	59 (7.4)	52 (6.4)	46 (5.8)	43 (5.2)		
10	169 (39.3)	117 (30.3)	81 (26.1)	65 (21.2)	62 (17.3)		
50	50 790 (44.5) 451 (42.4) 361 (32.6) 265 (28.7) 209 (27.0)						
100	1526 (47.3)	875 (45.2)	667 (35.7)	498 (31.7)	384 (30.5)		
Note:		1					

• Results may differ significantly from those listed here

Table 8.6. Kernel/Configuration - AS/400 S30-2257 and Network Station 300							
	Kernel/Configuration Initialization Time in Seconds						
	(Average CPU Utilization in %)						
		AS/400 Model S	S30-2257 (V4R1)				
	IBM	1 Network Station S	eries 300 (Releases 1	-2.5)			
		2629 16Mb Toke	en-Ring LAN IOP				
		8KB MTU/Frame	Sizes, 6 TFTP Jobs				
		Vary TFT	P Block Size				
# NS	512	1024	2048	4096	8192		
1	96 (1.8)	41 (4.3)	33 (4.5)	30 (3.7)	29 (3.6)		
10	182 (14.4)	73 (16.4)	56 (12.7)	52 (8.9)	39 (8.5)		
50	50 735 (18.9) 279 (24.9) 201 (20.2) 146 (17.5) 127 (15.3)						
100	1382 (20.2)	513 (27.7)	357 (23.2)	272 (20.0)	244 (16.6)		
Note:							

• Results may differ significantly from those listed here

Table 8.7. Kernel/Configuration - AS/400 400-2132 and Network Station 300						
	Kerne	el/Configuration Init	ialization Time in S	econds		
		(Average CPU l	U tilization in %)			
		AS/400 Model 4	00-2132 (V4R1)			
	IBN	I Network Station So	eries 300 (Releases 1	-2.5)		
		2629 10Mb Eth	ernet LAN IOP			
		6 TFT	'P Jobs			
		Vary TFTI	P Block Size			
# NS	512	1024	2048	4096	8192	
1	76 (35.6)	53 (26.3)	45 (19.8)	39 (17.7)	34 (15.5)	
10	10 280 (90.2) 167 (82.0) 110 (72.6) 83 (63.7) 67 (55.7)					
50	1311 (97.5)	745 (93.8)	467 (88.6)	321 (82.1)	277 (69.3)	
100	2591 (97.8)	1466 (96.9)	895 (93.4)	623 (86.7)	540 (73.1)	
Note						

Note:

• Results may differ significantly from those listed here

Table 8.8. Kernel/Configuration - AS/400 400-2132 and Network Station 300						
	Kerne	el/Configuration Init	ialization Time in Se	econds		
		(Average CPU l	Utilization in %)			
		AS/400 Model 4	00-2132 (V4R2)			
	II	BM Network Station	Series 300 (Release	3)		
		2629 10Mb Eth	ernet LAN IOP			
		6 TFT	P Jobs			
		Vary TFTI	P Block Size			
# NS	512	1024	2048	4096	8192	
1	64 (33.1)	51 (22.9)	46 (15.9)	42 (13.4)	40 (10.6)	
10	345 (82.8)	200 (75.0)	122 (62.4)	89 (51.1)	72 (48.4)	
50	50 1831 (93.3) 1109 (89.9) 533 (85.6) 334 (81.7) 274 (67.6)					
100 3678 (93.9) 1985 (90.1) 1111 (88.6) 660 (85.2) 543 (72.1)						
Note:						
Results may diffe	r significantly from th	hose listed here				

Table 8.9. Kernel/	Table 8.9. Kernel/Configuration - AS/400 400-2132 and Network Station 300					
	Kernel/Configuration Initialization Time in Seconds					
		(Average CPU	Utilization in %)			
		AS/400 Model 4	400-2132 (V4R2)			
	II	3M Network Station	Series 300 (Release	3)		
	Al	l NSs Attached to a	Single Twinax Adap	ter		
		6 TFT	'P Jobs			
	Vary Twinax Ada	pter Type, Subnet B	roadcast Option and	d TFTP Block Size		
# NS	6050 w/o	6180 with	6180 w/o	6180 with	6180 w/o	
	Subnet 8K	Subnet 1K	Subnet 1K	Subnet 8K	Subnet 8K	
	TFTP	TFTP	TFTP	TFTP	TFTP	
1	107 (8.5)	114 (22.0)	116 (22.1)	87 (14.5)	82 (13.3)	
2	173	155	133	90	85	
3	225	165	154	106	98	
4	275	168	159	121	116	
5	325	186	178	142	139	
6	388	201	199	155	157	
7	446 (16.4)	225 (33.6)	221 (70.0)	171 (28.8)	162 (37.5)	
Note:						
Results may diffe	r significantly from th	nose listed here				

Note that subnet broadcast uses less AS/400 CPU. However, as discussed above, each twinax device on the subnet will get their own copy of the broadcast data, even if they didn't request it, which would mean unwanted data on the twinax cable. In general, customers should not use twinax subnet broadcast. (Subnet broadcast should be used on LANs.)

In Table 8.9, the Network Stations were all chained to a single cable port. For the 6180 adapter, faster times could be obtained if the Network Stations were balanced across cable ports, that is, half on ports 0-3 and the other half on ports 4-7. For example, in the table above, 6 Network Stations with an 8K TFTP block size, without subnet broadcast, booted in 157 seconds. If they had been balanced, 3 on port 0 and 2 on port 4, the initialization time would have been 130 seconds, 17% faster.

If a Network Station has multiple paths, with the same network address, to an AS/400 (e.g. two IOPs that each have a path to the Network Station), unexpected results may occur. Whenever the AS/400 gets a request from a Network Station, it uses the default path to get back to the requesting station. The return route (and any subsequent request/replies) may be different from the original request. This implies that there is no value to add a second IOP with the same network address to gain additional TFTP performance.

TFTP jobs are assigned first come, first serve. There is no mechanism to allocate a TFTP job to a particular IOP. This implies that it is possible for Network Stations attached to one network to monopolize all the TFTP jobs until completion of the kernel download. Other IBM Network Stations maybe starved until a TFTP job is available.

3. Login

Login is just that - the user enters his/her user-ID and password and then the desktop appears.

The load times can be found in the table below.

4. Application Load

Applications are loaded when their respective desktop buttons are selected. Load times vary by AS/400 machines size.

Getting to a 5250 sign-on can require two steps - from the menu bar, select the 5250 button to get to the host name window and then enter the desired host name to get to the 5250 sign-on window. Most administrators use the Network Station Manager to configure for direct menu bar to 5250 sign-on.

Getting to the browser is a single step - from the menu bar, select the browser button to get to the browser.

Examples of load times can be found in the table below.

Table 8.10. Application Load Times							
		Load Times in Seconds					
	AS/400 Mo	odel 150-2270 and 510-2	144 (V3R7)				
	IBM Network St	tation Series 100 and 30	0 (Releases 1-2.5)				
	2619	16Mb Token-Ring LAN	IOP				
2270 100 2270 300 2144 100 2144 300							
Login to desktop	10	10	15	11			
5250 select to host	9	6	10	7			
name							
Host name to 5250	6	6	12	11			
login	login						
Browser select to	33	16	41	22			
browser							
Note:							
• Results may differ significantly from those listed here							

• Results may unter significantly from those listed h

Table 8.11. Application Le	oad Times						
Load Times in Seconds AS/400 Model 400-2132 (V4R2) IBM Network Station Series 300 (Release 3) eSuite is IBM Network Station 1000 (Release 3)							
	Twinax or Ethernet Adapter 6050 6180 2620 2620 DBCS*						
Login to desktop	30	27	18	23			
5250 select to host name	57	33	10	19			
Host name to 5250 login	15	21	12	14			
Browser select to browser	169	131	41	52			
eSuite select to eSuite			175				
Note:	Note:						

• Results may differ significantly from those listed here

• *DBCS support includes Japanese, Korean, simplified Chinese, and Traditional Chinese

Another example of subnet broadcast: Assume 100 Series 300 Network Stations attached to an AS/400 V4R2 2132 via a single 10Mb ethernet segment. Assume the electricity on all 100 Networks Stations goes out and a while later comes back on. Assume the Network Stations all

have the same memory size (e.g. 32MB) and identical monitors attached. It would be possible for all 100 to be at the Login window in 280 seconds (less than 5 minutes). The 280 seconds comes from: 21 seconds for hardware test, 30 seconds to load the kernel, and 229 seconds to load configuration files.

8.3 AS/400 5250 Applications

The Network Station user should see 5250 applications almost exactly as with NPT or PC terminals. However, the load on the AS/400 may be different. Network Stations use the AS/400 TCP/IP Telnet path. Telnet consumes 27% more CPU time per transaction than an NPT attached to a local twinax for a typical commercial workload. This yields a 20% capacity reduction over a twinax attached NPT. For comparison, a Client/Access PC using 5250 over SNA, when using the same workload, consumes 10% more CPU time per transaction than a local twinax attached NPT.

The implication is that customers migrating from local twinax attached NPTs to LAN attached Network Stations will probably use more CPU to run the same 5250 applications. Customers migrating from LAN attached SNA Client/Access PCs will also probably use more CPU. Customers migrating from LAN attached TCP Client/Access PCs should need no additional CPU capacity to run their 5250 applications.

8.4 Browser

In general, the Series 100, 300 and 1000 all perform equally well. Their performance should be comparable to that seen on a PC.

It is important that either socks or proxy are configured, but not both. Poor performance is seen when both are used.

Disk caching should never be used.

8.5 Java Virtual Machine Applets/Applications

Java is still evolving. As such, its use on a Network Station is also evolving. The Series 100 clearly should not be used for Java. The Series 300, while twice as fast as the 100, can be used for very limited Java applets. The Series 1000 is for Java; however, since Java has varied uses, customers are encouraged to test their Java applications on the Series 1000 before putting them into production.

8.6 The AS/400 as a Router

The AS/400 is a router (data passes though it) when twinax attached Network Stations send/receive data from the internet or other servers. At this time, limited performance data is available. The following two tables show results when data is read from an NT server through an AS/400 to a Network Station.

Table 8.12. LAN to LAN Throughput

LAN to LAN Throughput AS/400 Model 400-2132 (V4R2) IBM Network Station Series 300 (Release 3) via 10Mb Eth to AS/400 300MHz PC NT server via 16Mb TR to AS/400 2629 LAN IOPs, 15MB of Data, 8K TFTP Block

	/		
# NS	Time (sec)	AS/400 Util (%)	AS/400 Throughput (Kb/s)
1	44	11.2	340.9
2	48	16.9	625.0
3	57	18.1	789.5
4	71	25.7	845.1
5	90	24.4	833.3
10	158	29.9	949.4
15	232	34.9	969.8

Table 8.13. LAN to Twinax Throughput				
LAN to Twinax Throughput				
	AS/400 Model 4	400-2132 (V4R2)		
IB	M Network Station Series 300	(Release 3) via Twinax to AS/	400	
	300MHz PC NT server	via 16Mb TR to AS/400		
2629	LAN IOPs, 6180 Twinax Ada	pter, 2MB of Data, 8K TFTP	Block	
			AS/400 Throughput	
# NS	Time (sec)	AS/400 Util (%)	(Kb/s)	
1	33	9.9	70.1	
2	48	9.9	96.4	
3	109	10.3	63.7	
4	127	10.5	72.9	
5	150	11.1	77.1	
6	213	11.0	65.2	

8.7 Conclusions

The IBM Network Station provides for an excellent working environment.

- In general, the Network Station 1000 performs better than the 300 which performs better than the 100
- Initialization
 - The Network Station Series 1000 initialization time is about the same as the 300, except for hardware test, where the 1000 is faster. The 300 is faster than the 100.
 - If possible, customers should consider a boot server for each ring or ethernet.
 - For Releases 1-2.5, customers should use BOOTP or DHCP and not NVRAM. BOOTP and DHCP are faster and more reliable. For Release 3, all three initialization mechanisms are equal in reliability and performance. BOOTP is slightly (1-2 seconds) faster than DHCP.
 - The time to initialize Network Stations depends on many variables, such as size of AS/400, TFTP block size, number of attached IBM Network Stations, LAN utilization, CPU utilization, etc. Customers will need to evaluate their own needs. It is recommended that customers go slow in

building their Network Station solutions.

- Initialization time varies from AS/400 model to AS/400 model. In general, the faster the model, the better the performance. On faster models, the bottleneck is the LAN IOP and, on slower models, the bottleneck is CPU and LAN IOP. The 2629 LAN IOP provides better performance than the 2619.
- ◆ 10Mb Ethernet, 100Mb Ethernet and 16Mb token ring are about equal.
- During initialization, CPU utilization can be quite high, especially on the smaller AS/400s, which will impact other jobs. In addition, TFTP requires more CPU than RFS.
- Subnet broadcast can significantly reduce LAN traffic and AS/400 CPU utilization. Subnet broadcast is available with AS/400 V4R2 and Network Station Release 3. If possible, it is highly recommended that subnet broadcast be used. In general, subnet broadcast is not advisable with twinax, except as discussed earlier.
- The network administer should configure TCP/IP, LAN frame size and TFTP block size for best performance. In general, the larger the size, the better the performance.
- For twinax, the 6180 adapter is significantly faster than the 6050. The 6180 is about equal to a 4Mb token ring.
- There is no value to add a second IOP, with the same network address, to a LAN to get better initialization performance, since TFTP will select the path to be used. All Network Stations, from the same network, will use the TFTP selected path.
- It is best to configure 6 TFTP jobs per LAN that has attached Network Stations. However, for systems that have multiple LANs, since there is no way, at this time, to dedicate a TFTP job to a particular LAN, initialization may not perform as well as desired.
- In general V4R2 provides better performance than V4R1, which provides better performance than V3R7. V4R2 contains TCP/IP and IOP LAN enhancements. In some cases, customers will see substantial improvements in kernel/configuration initialization. These improvements, in general, will be visible when a single Network Station is initialized with a small TFTP block size. V4R2 contains RFS enhancements. V4R3 and V4R4 performance is the same as V4R2.
- Release 3 boots about as fast as previous releases, even though more data and function are sent. Much of the data sent is compressed.
- Switches, routers and gateways can cause problems. It is best to have a network administrator.
- For 6180 twinax attached Network Stations, best performance is obtained if all Express Datastream enabled devices are on the same cable, excluding older, non-Express capable devices.
- When Express devices are attached to a single workstation controller, best performance is obtained by load balancing those devices. That is, half the devices should be connected to cable ports 0-3, and the other half should be connected to ports 4-7.

- 5250 application performance on the AS/400
 - ♦ In general, the 100, 300 and 1000 all perform equally.
 - Customers migrating from LAN attached SNA Client/Access PCs will probably use more CPU (about 17%) to run the same 5250 applications.
 - Customers migrating from LAN attached TCP/IP Client/Access PCs will use about the same CPU to run the same 5250 applications.
 - Customers migrating from local twinax attached NPTs to IBM Network Stations will probably use more CPU (about 27%) to run the same 5250 applications.
- Browsers
 - ✤ In general, the 100, 300 and 1000 all perform equally.
 - Poor performance is obtained when both socks and proxy are configured. Only one should be used.
 - ✤ Never use disk caching.
- Java Virtual Machine
 - The Series 100 should not be used for Java.
 - The Series 300 can be used for limited, lightweight Java
 - The Series 1000 is for Java; however, since Java hasn't fully matured and can be used for many, varied applications, customers should insure that their Java application and the 1000 are compatible.
- AS/400 as a Router
 - Limited performance data is available. A model 400-2132 is able to route about 970kb/s from one LAN to another and about 75Kb/s from a LAN to twinax.

Chapter 9. AS/400 File Serving Performance

This chapter will focus on AS/400 File Serving Performance.

0In the V4R4 update of this chapter, the content that pertained to previous AS/400 releases has been removed. The information is still available in the V4R3 Performance Capabilities Reference.

9.1 AS/400 File Serving Performance

In V4R4, performance improvements were made to the Integrated File System (IFS). The V4R4 enhancements affect the Root, QOpensys, and User-Defined File Systems. The other file systems (Qsys,QDLS, Qopt...) will function at the same level of performance as the previous release.

The Pre-bring buffering schemes were improved, along with other changes, resulting in WRITES being upto 2x faster. As the number of files being accessed and the size of these files increase, the degree of improvement decreases to the point where there might not be any noticeable change in performance.

9.2 AS/400 NetServer File Serving Performance

AS/400 NetServer supports the Server Message Block (SMB) protocol through the use of Transmission Control Protocol/Internet Protocol (TCP/IP) on AS/400. This communication allows clients to access AS/400 shared directory paths and shared output queues. PC clients on the network utilize the file and print-sharing functions that are included in their operating systems. You can configure AS/400 NetServer properties and the properties of AS/400 NetServer file shares and print shares with Operations Navigator.

Clients can use AS/400 NetServer support to install Client Access from the AS/400 since the clients use function that is included in their operating system. See <u>http://www.as400.ibm.com/netserver/</u> for additional information on AS/400 NetServer.

V4R4 Note: Tables and Figures were added to reflect the V4R4 changes.

In V4R4, enhancements were made to AS/400 NetServer . The V4R4 enhancements affect the dally timer delay cycle, and TCP Send and Receive buffer sizes. When a user migrates from Network Drives with Client Access to AS/400 NetServer on V4R4, they can expect to see an improvement in performance.

V4R4 AS/400 NetServer Performance

Server

 AS/400 Invader Model 2292/170 V4R4
 1 GB Memory, 935,330 MB for Base Pool, 9- 8GB Disk Drives 16Mbps Token Ring LAN - 2724 IOP

Clients

- Pentium 133MHz, 32 MB Memory, 1.2GB IDE Disk Drive, 16Mbps Token Ring PCI Adapter Client Access Express for Windows NT (This product was installed for comparison purposes only to provide a comparable PC environment. Connection of the network drive to AS/400 NetServer was done using the function provided with Windows NT.) Windows NT Workstation 4.0
- Pentium 133MHz, 32 MB Memory, 1.2GB IDE Disk Drive, 16Mbps Token Ring PCI Adapter Client Access for Windows NT V3R1M2 Windows NT Workstation 4.0

Workload

200MB File Transfer in both directions (Upload and Download using DOS copy command)

Measurement Results:



Conclusion/Explanations:

From the chart above in the Measurement Results section, it is evident that when customers migrate from Network Drives with Client Access to AS/400 NetServer on V4R4 they can expect to see an increase in peformance. It is also clear from the chart that AS/400 NetServer achieves nearly the same performance in both directions when transferring large files

Chapter 10. DB2/400 Client/Server and Remote Access Performance

With the announcement of the AS/400 Advanced Servers using PowerPC technology, IBM has again provided a clear direction for customers who are moving to a client/server environment. With V4R1, overall system performance of the high-end AS/400 server model S40 12-way increased significantly over the existing high-end Model 53S, giving outstanding growth and improved price/performance. For customers who have a mixed environment (a combination of fixed function workstations and PC's), then AS/400 Advanced System models using PowerPC technology provide significant system performance growth and improved price/performance (see Chapter 2, "AS/400 System Capacities and CPW" for relative performance of these models).

When using client/server technology, it is important to consider the impact of the various client and server components, and their effect on performance. There are different ways of implementing client/server applications. In this chapter, guidelines are provided for a number of common implementation strategies, to help understand the impact of the performance of the client system and the server system using database serving workloads.

With the introduction of Java on the AS/400, database access and client/server development have changed to an even more open environment. The AS/400 Toolbox for Java provides a JDBC driver that conforms to the JDBC specification published by Sun Microsystems. JDBC enables application developers to write portable applets and applications that access relational database information.

Client Access/400 contains an OLE/DB driver for the AS/400 in V4R1. This driver allows developers to easily and quickly develop client/server applications for the AS/400.

ODBC and JDBC will continue to be strongly supported by IBM as an *open* way of connecting to DB2/400.

In **V4R4**, NetServer provides file serving performance comparable with Network Drive when comparing different types of AS/400 with the same CPW value. Please refer to section measurement results section of the section entitled *Server Challenge Benchmark(SCB)* for more information.

In general, V4R2 database serving performance will be equivalent to V3R2 when comparing AS/400 models with the same CPW value. The new models using PowerPC technology in this case provide significant system performance growth and improved price/performance. See "Server Challenge Benchmark (SCB)" 132 for more information on AS/400 release comparisons.

Use the information provided in *AS/400 Performance Capabilities Reference (V3R2), ZC41-8166*, Chapter 7, "DB2/400 Client/Server and Remote Access Performance Information", as a guide for V3R2 performance. In addition, refer to "Related Publications/Documents" at the beginning of this document on how to access a presentation that covers *AS/400 Versus Microsoft's SNA Server Gateway*.

10.1 Client Performance Comparisons

Under different client server implementations, PC hardware configuration plays an important role in overall performance. In general, a faster PC CPU will improve performance, but there are other issues which should be taken into account such as disk drive performance, main storage, memory cache etc. If an

acceptable relative performance value was 2.0 (or 2X slower than a '486 @ 66 MHZ), then the PS/2 model 80 would prevent this environment from achieving the required response time criteria.

For applications where 50% or more of the application response time contribution is on the PC client, such as a query download or an OLTP application, better performance can be achieved by focusing on the client performance and selecting a faster 486, Pentium (**), Pentium Pro (**), or Pentium II (**) processor.

It is also important to note that client memory can be a significant component of response time as well. For the Windows 3.1 client, most database serving operations will perform acceptably if the client has at least 8 MB of memory. The Windows 95 client performs acceptably with 32MB of memory, and the Windows NT client performs acceptably with 32MB of memory. Client/Server operations that operate on a large amount of data will usually perform better if the client has more than the amount of memory suggested above.

When most of the response time contribution is on the AS/400 such as a complex query, or when processing an OLTP stored procedure, greater performance improvements may be achieved by optimizing the AS/400 application or upgrading AS/400 hardware. For example, it may be possible to create a logical view for a query which is frequently executed.

10.2 AS/400 Toolbox for Java

The AS/400 Toolbox for Java is a set of enablers that supports an internet programming model. It provides familiar client/server programming interfaces for use by Java applets and applications. The toolbox does not require additional client support over and above what is provided by the Java Virtual Machine and JDK.

The toolbox provides support similar to functions available when using the Client Access/400 APIs. It uses sockets connections to the existing OS/400 servers as the access mechanism for the AS/400 system. Each server runs in a separate job on the AS/400 system and sends and receives architected data streams on a socket connection.

The AS/400 Toolbox for Java is delivered as a Java package that works with existing servers to provide an internet-enabled interface to access and update AS/400 data and resources.

The base API package contains a set of Java classes that represent AS/400 data and resources. The classes do not have an end-user interface but simply move data back and forth between the client program and an AS/400 system, under the control of the client Java program.

For more information on the AS/400 Toolbox for Java see the Redbook *Accessing the AS/400 System with Java* Document Number SG24-2152-00.

JDBC Driver

The JDBC driver that is included in the AS/400 Toolbox for Java allows database access to the AS/400 using APIs that are similar to ODBC. This JDBC driver talks to the same server job on the AS/400 as the ODBC driver included with Client Access/400. Many of the options for the Client Access/400 ODBC driver are therefore included in the JDBC driver. Also, any of the database and communication tuning for ODBC and Client Access/400 can be used for JDBC and the Toolbox.

JDBC allows SQL statements to be sent to the AS/400 system for execution. If an SQL statement is run more than one time, use a PreparedStatement object to execute the statement. A PreparedStatement compiles the SQL once, so that subsequent executions run quickly. If a plain Statement object is used, the SQL must be compiled and run every time it is executed. Use Extended Dynamic support; it caches the SQL statements in SQL packages on the AS/400 system. Also turn on package cache; it caches SQL statements in memory.

Do not use a PreparedStatement object if an SQL statement is run only one time. Compiling and running a statement at the same time has less overhead than compiling the statement and running it in two separate operations.

Consider using JDBC stored procedures. In a client/server environment, stored procedures can help reduce communication I/Os and thus help improve response time.

Use a just-in-time (JIT) compiler for your Java execution environment if possible. The latest JIT technology allows Java programs to perform almost as well as native code written in C or C++. Users of the AS/400 Toolbox for Java can expect the JDBC driver to perform almost as well as a C++ application using ODBC for OLTP types of applications. JDBC applications that download larger amounts of data will perform slower than a comparable C++ application in ODBC because of the object orientated design of the JDBC driver.

There are many properties that can be specified on the JDBC URL or in the JDBC properties object. Several of these properties can significantly affect the performance of a JDBC client/server application and should be utilized where possible. The properties control record blocking, package caching, and extended dynamic support. See the JDBC driver documentation for details on setting these properties. Most of the properties have close parallels in the ODBC driver. Tuning advice from ODBC can be used for JDBC when setting these values.

Record Level Access

AS/400 physical files can be accessed a record at a time using the public interface of these classes. Files and members can be created, read, deleted, and updated. The record format can be defined by the programmer at application development time, or can be retrieved at runtime by the AS/400 Toolbox for Java support. These classes use the DDM server to access the AS/400 system. To use the host DDM server through a TCP/IP interface, some special PTFs are required. Check the AS/400 Toolbox for Java documentation for the latest PTF numbers and set up instructions.

Record Level Access can offer better performance than JDBC for applications that need to process AS/400 database data one record at a time. Record access does not go through the SQL query processing that JDBC must go through to process data, therefore it can retrieve a single record quicker than JDBC. However, if complex computations or large sets of records are processed, JDBC may be a better solution. Use JDBC when an SQL statement can be built that does most of the work on the server, or when you want to limit which fields in a record get transfered to the client. The current JDBC specification does not have a mechanisim to insert multiple records at a time (e.g., ODBC's Blocked Insert). Therefore, use the support in Record Level Access to write multiple records at once.

There are several issues that should be considered when using the Record Level Access support in the Toolbox. First, when accessing a file, if the program is going to use the file multiple times, the file should

be left open inbetween operations to avoid the extra processing due to open and close. Second, the block size that is specified on the open method should be selected acording to the type of access in the file. Block size is the number of records to download to the client when reading. If multiple records that are relatively close together are going to be retrieved then a block size that can transmit all of the records at once is prefered. However, do not select a block size that will cause a large delay when downloading (e.g., a 1MB download). If the type of access for the file is random, and the records retrieved are not close together in the file, a block size of 1 is prefered. Third, when downloading an entire file that is relatively small, use the readAll method. This method is considerably faster when reading small files, large files (e.g. >1MB) may encounter "Out of memory" errors, beacuse the entire file is placed and translated in the client's memory.

10.3 Client Access/400

Client Access for Windows 95/NT (XD1)

In V3R1M2 Client Access/400 added support for Windows NT. The product Client Access for Windows 95 has been replaced with Client Access for Windows 95/NT.

Client Access for Windows 95/NT communications support provides excellent performance with native TCP/IP connectivity. Additionally, its native APPC support is equal to or better than most other 32-bit APPC Router implementations available for Windows 95/NT. For optimum client/server application response time performance, we recommend that customers use TCP/IP connectivity.

There are several migration scenarios involved with the AS/400 Client Access for Windows 95/NT client, and each one has different performance characteristics:

- 1. Customers using APPC connectivity on the AS/400 Client Access for Windows 3.1 client or DOS Extended client may experience a client response time degradation when migrating to AS/400 Client Access for Windows 95/NT *APPC* connectivity. AS/400 server capacity is unchanged.
- Customers using APPC connectivity on the AS/400 Client Access for Windows 3.1 client or DOS Extended client will in most cases experience equivalent client response times when migrating to native *TCP/IP* connectivity on the Client Access for Windows 95/NT client. AS/400 server capacity may be somewhat reduced.
- 3. Customers using AnyNet (APPC over TCP/IP) connectivity on the AS/400 Client Access Windows 3.1 client will, in most cases, experience improved client response times when migrating to native TCP/IP connectivity on the Client Access for Windows 95/NT client. Native TCP/IP connectivity is available to all Client Access applications or to applications written to Client Access for Windows 95/NT 32-bit Application Programming Interfaces (API's). Applications written to Windows 3.1 Client Access API's, will run over SNA or, in addition, over TCP/IP by using AnyNet (APPC over TCP/IP). AS/400 server capacity will be improved in most cases.
- 4. Customers migrating from Client Access for Windows 95 to the new product, Client Access for Windows 95/NT, should experience equivalent client response times. Users changing client operating systems from Windows 95 to Windows NT and using the new Client Access/400 may notice a slight improvement in response time for some operations.

5. Slower clients ('486 66MHZ or less) running SNA may experience a performance degradation when migrating to the Client Access for Windows 95/NT from Client Access for Windows 95. The level of degradation depends on the environment. Typical operations using SNA with slower clients should be no more than 10-50% slower than Client Access for Windows 95/NT. Faster clients may experience very little, if any, degradation.

Due to the robust nature of the SNA router connectivity in the Client Access Windows 95/NT client, startup response times may be significant, especially on PC's with lower amounts of memory. The Windows 95 client performs best on a 486 66 MHZ client or greater system with at least 16 MB of memory. The Windows NT client performs best on a Pentium 133 MHZ or greater system with at least 32 MB of memory. Refer to Client Access information for the minimum system sizes supported.

After the first connection has been established to an AS/400, subsequent connections to the system or any other system are much faster. Because of this, it is best to leave open connections up, if possible, and not reconnect each time you use an application running over this connectivity.

ODBC users should be aware that a parameter setting has changed with the new Client Access for Windows 95/NT. A PREFETCH setting was added with the default set to OFF. Previously, PREFETCH was built-in and always set to ON. This setting was added because PREFETCH is not supported for applications that use SQLExtendedFetch commands. Some users may see a performance degradation if PREFETCH is not enabled. Users that use SQLExtendedFetch should leave this setting OFF.

OLE/DB and ADO Data Access (Project Lightning)

The OLE/DB driver that has been added to the base Client Access/400 for Windows 95/NT allows database access to the AS/400. Developers can write applications that use the OLE/DB driver to access the AS/400 database through DDM Record Level Access, SQL, Stored Procedures, etc. These interfaces can be easily programmed through the ADO layer that most current development environments support (e.g., Visual Basic, Delphi, etc).

The current ADO specification does not support record blocking; therefore, downloading a large table through record level access may take longer than other methods (e.g., ODBC which has record blocking) and consume more network resources, since each record is transmitted as one communication. Look for future versions of the ADO specification to contain record blocking.

The SQL support in the OLE/DB driver uses the same server program as the ODBC driver. This means that developers can use some of the same techniques and ideas from the Client Access ODBC driver for the OLE/DB driver. One important performance improvement that developers can use, is implement prepared statments if an SQL statement is to be executed more than once. Also when executing a prepared statment set the third parameter to be -1, otherwise ADO assumes this is a new statement and discards the previously prepared statement.

Open Data Base Connectivity - ODBC()**

In V3R1, the ODBC APIs were significantly enhanced in terms of function and performance compared to the original version in V2R3 (see *AS/400 Performance Capabilities Reference (V3R2), ZC41-8166*, Chapter 7, for more details). ODBC support for the Windows (**) 3.1/95/NT clients in Client Access/400 provides superior performance over the original Remote SQL support in Client Access/400 and its predecessor PC Support/400.

ODBC is a set of API's (Application Programming Interfaces) which provide clients with an open interface to any ODBC supported database. AS/400 supports ODBC with the Windows 3.1, Windows 95, Windows NT and OS/2 client support in Client Access/400. Customers can purchase ODBC drivers to connect to their system(s) and either write applications which utilize the ODBC APIs or purchase an existing application which utilizes the ODBC APIs.

Client/Server 4GL and Middleware

Many users build client/server applications using client toolkits such as C/S fourth generation languages (4GLs). Most of the new 4GL tools use "middleware" or interface code to connect to a server. This middleware usually consists of one or more DLLs used to connect to a given server. The middleware converts the client's request into commands and data which the server can understand and converts the server's response into commands and data which the client can understand. Often the middleware is written by the toolkit provider to interface to a given server or to a standard server API set.

Examples of 4GL toolkits are GUPTA's SQLWindows**, PowerSoft's Powerbuilder**, Microsoft's VisualBasic**, and Visual C ++. Examples of middleware standards are Microsoft's ODBC standard and IBM's DRDA standard.

Because the user is often isolated from the APIs and the middleware manages the database access method, it is important to build applications using tools that optimize for performance. In many cases, tools that are built for "openness" for many servers tend to be the worst performers because they are built to the least common denominators. The AS/400 supports many features that enhance performance. Ensure that your toolkit has support for functions like stored procedures and blocked insert. If not, ensure that there is a mechanism to write directly to the CA/400 API set for the best performance. For more information on client/server application development tools, see *AS/400 Client/Server Performance Using Application Development Tools, (SG24-4731)*.

Client Toolkit ODBC Performance

Many performance problems with client development toolkits are due to the client tool creating inefficient database access requests. For example, a simple database transaction that should result in minimal interaction with the server can generate hundreds of unnecessary ODBC requests and responses. By choosing high performance toolkits and with planning and tuning, these problems can be avoided.

Tools are available to diagnose and debug problems with client toolkits and applications. Use tools such as ODBCSpy or ODBC Trace (available through the ODBC Driver Manager) to verify the efficiency of the SQL and ODBC calls that are generated. Also, the toolkits themselves often have tools to trace their server access methods.

Client/Server Online Transaction Processing (OLTP)

OLTP applications are typically designed for business computing. An OLTP transaction usually consists of several database operations and related computations. Performance requirements for this type of transaction are usually stringent. In the client/server environment an OLTP transaction typically consists of several requests/responses between the client and the server, resulting in a small to moderate amount of data transferred to the client. Because these transactions tend to be repetitive, they are good candidates for application serving (such as remote procedures or distributed processing). It is especially important to

avoid unnecessary overhead in processing these repetitive transactions (such as PREPARE operations). Use of parameter markers, stored procedures, triggers, and Extended Dynamic (package) support is recommended to improve the performance of this class of queries.

Server Challenge Benchmark (SCB)

SCB Overview: The Server Challenge Benchmark (SCB) is a set of three individual workloads developed by IBM to compare the AS/400 against the competition:

- 1. Transaction-based Component Workload
- 2. Decision Support Component Workload
- 3. File Server Component Workload

Each of these component workloads were built to use client/server characteristics. For a detailed description of SCB refer to the IBM white paper titled "AS/400 Client Server Performance Benchmark Guide" available on HONE.

This application is typical of a commercial client/server application where many small transactions are being continuously processed. It is not representative of a decision support workload nor that of a file server workload, which are the other two components of the SCB. Additional information on file server performance can be found in Chapter 9, "AS/400 File Serving Performance".

SCB with Windows 3.1/95/NT Client:

Measurement Configuration:

AS/400 170 - 2292 - dedicated V4R4 1GB Memory 9-8GB Disk Drives (72 GB) 16Mbps Token Ring Lan - 2724 IOP

Clients:

- Win NT client Pentium-133 MHZ 32MB memory Windows NT 4.0 1.2GB IDE Disk Drive Client Access Express for Windows
- Win NT client Pentium-133 MHZ 32MB memory Windows NT 4.0 1.2GB IDE Disk Drive Client Access for Windows 95/NT V3R1M2
- 16 Mbps TR-LAN SCB Transaction-based Workload - implemented with Visual C++ One client - Average transaction response time - 100 transactions (All response times are expressed in seconds)

V4R4 SCB Performance

In V4R4, when migrating from Network Drives to NetServer, the customer can expect to see quivaence or sightly better performance...

Client/Server Query and Decision Support

Query and Decision Support (also known as ad-hoc queries) are database operations typically done throughout the day in many businesses. These operations are longer-running server-intensive database operations which are usually read-only. Unlike OLTP, they are seldom done frequently and although throughput may not be a critical factor in performance, response time surely is. These queries usually result in minimal interactions between the server and the client. Although they may return a large number of rows of information, typically ad-hoc queries return few rows. Because these queries are server-intensive, they are good candidates for database serving. The response time of these remote (c/s) queries are not significantly slower (typically 5% - 30%) than local (host terminal) queries. However, since the server performance load can be large, the user may benefit from moving the execution of these queries off-shift.

In cases where the query executed is CPU-intensive and very few records are returned to the client, the response time of the query is typically very close to running the same query interactively on the server.

ODBC Query Using Re-use Pre-started Jobs

V3R7 OS/400 provides a new option to re-use pre-started jobs for the QSERVER subsystem. This allows pre-started jobs to "recycle" jobs that have previously ended. As a result, CPU time to start-up jobs is reduced with the re-use value set to greater than 1.

This value can be changed with the AS/400 command CHGPJE. Use the QSERVER subsystem with the job QZDAINIT and subsystem QIWS for APPC jobs or QZDASOINIT for TCP/IP or IPX jobs. Press F10 for additional parameters then page down to the parameter "MAXIMUM NUMBER OF USES". Enter the new value for the maximum job re-uses. The V3R7 default value is 200.

Query Download/Upload (Database file transfer)

Download/Upload queries represent a set of queries that either fetch a significant number of rows from DB2/400 tables (or files) or insert a significant number of rows into DB2/400 tables or files. Because of the number of rows processed, there is a significant amount of processing that occurs on the client. Many times, significant performance gains may be realized by running these types of queries on a faster client processor.

Query Download Comparisons

This section compares the performance of downloading a significant number of records from an DB2/400 database file into a client application using various CA/400 APIs. To obtain the download comparisons, a client application was developed that fetches about 1.4 MBs of data using different APIs. All fetched character data was converted from EBCDIC to ASCII automatically by CA/400 functions. No further processing was performed on the data retrieved.

Measurement Configuration:

The following configuration was used to perform the query download measurements:

```
AS/400 Server 50S-2120 - dedicated
256 MB Memory
2-6606, 2-6605 DASD (5.99 GB)
ValuePoint clients - '486-66 MHZ - 32MB memory - CA/400 for Windows 3.1 V3R1M1
2048 byte frame size
16 Mbps TR-LAN
```

The download operations were done using standard PC-based query tools. The tool was changed to display the first resulting rows or to simply indicate completion of the test. For query download tests, the ODBC.INI file on the client was changed to vary the Record Blocking size setting. The following queries were performed to download records from the DB2/400 database files:

Table 10.1.	CA/400 Windows 3.1 Query Download Tests				
		Row	Columns	Rows	Bytes
Query	SQL Statement	Size	Per Row	Fetched	Fetched
LBR	SELECT * FROM DBITRK/LBRSTATS	118	10	12,000	1,416,000
TRK	SELECT * FROM DBITRK/TRKOPRNS	314	42	4,418	1,387,252
	WHERE TRKITEMN				
	<'ITE02210MNUMBR'				

Measurement Results:

The following table shows the ODBC.INI Record Blocking size setting, the overall download rates in megabytes per hour, the overall response times in seconds, and the AS/400 CPU seconds consumed for the queries listed above. The queries were implemented using Client Access/400 ODBC.

Table 10.2. CA/400 Windows 3.1 Query Download Performance					
	Record	Transfer	Response Time	AS/400 CPU Consumed	
Query	Block Size	Rate (MB/hr)	(seconds)	(seconds)	
LBR - display rows	512K	296	17.2	0.63	
LBR - no display	512K	566	9.0	0.63	
LBR - no display	32K	520	9.8	0.69	
TRK - display rows	512K	271	18.4	0.78	
TRK - no display	512K	450	11.1	0.78	
TRK - no display	32K	427	11.7	0.86	
Note: All tests received all rows: "	'display" tests only dis	played the first set of ro	ows		

Query tools are available to provide the client with an easy, graphical way to access server databases. For performance reasons, many of these tools allow the user to limit the number of rows received. We ran one of the above queries and limited the rows using a query tool. This table shows the row limit, the ODBC.INI Record Blocking setting, and the overall response times in seconds.

Table 10.3. CA/400 Windows 3.1 Query Download Performance - Limited Rows					
Response					
		Record Block	Time		
Query	Row Limit	Size	(seconds)		
LBR	20	32K	1.4		

Query	Row Limit	Record Block Size	Response Time (seconds)					
LBR	20	512K	3.3					
TRK	20	32K	2.2					
TRK	20	512K	4.1					
Note: These tests only downloaded the f	irst 20 rows of data	Note: These tests only downloaded the first 20 rows of data						

For comparison, we measured an IFS file transfer (download) operation using one of the same database tables above. We converted the DBITRK/LBRSTATS table to a stream file and loaded it on the server in the IFS root directory. We then used Windows File Manager to copy the file from the server to the PC hard disk. We compared the time for downloading a new file with replacing an existing file. The following table shows the overall download rates in megabytes per hour, the overall response times in seconds, and the AS/400 CPU seconds consumed for the queries listed above.

Table 10.4. CA/400 Windows 3.1 File Download Performance					
		File Transfer	AS/400 CPU		
	File Transfer	Time	Consumed		
File Downloaded	Rate (MB/hr)	(seconds)	(seconds)		
DBITRK/LBRSTATS - new file	593	8.6	0.59		
DBITRK/LBRSTATS - replace file	614	8.3	0.51		
Note: These tests downloaded the file to	the PC hard disk				
Conclusions/Recommendations:

- 1. Only a small percentage of the total transfer time was due to the AS/400 CPU. Most of the time spent for large record download operations is in the client and communications time. Use fast clients for the best performance. Use fast communications adapters for higher throughput.
- 2. ODBC query download rates can be comparable to IFS file transfer rates.
- 3. For fastest retrieval times for an entire large database table, do not immediately format and display all the data retrieved. Instead, use client tools to manipulate and display the data after it has been entirely downloaded to the client.
- 4. When retrieving the entire database table, the recommended ODBC Record Blocking setting is 512 KB. Decreasing this size may cause slower performance. Memory-constrained clients may require a smaller block setting.
- 5. When using client tools to browse through the data, limit the query to display only the first screen of data. Fetch the next set of data when needed. Set the Record Blocking to 32K or less for fast retrieval of only a small number of rows from a large table.
- 6. As the number of columns to be retrieved increases, the retrieval rate decreases and response time increases.
- 7. The token-ring frame size used was 2K. Larger frame size settings may improve performance.

Query Upload Scenario

This section compares the performance of uploading a significant number of records into a DB2/400 database file from a Windows 3.1 application using various CA/400 ODBC APIs.

Measurement Configuration:

The following configuration was used to perform the query upload measurements:

```
AS/400 Server 50S-2120 - dedicated
256 MB Memory
2-6606, 2-6605 DASD (5.99 GB)
ValuePoint clients - '486-66 MHZ - 32MB memory - CA/400 for Windows 3.1 V3R1M1
2048 byte frame size
16 Mbps TR-LAN
```

The client application is written in C and utilizes CA/400 Windows 3.1 ODBC APIs to do single inserts and blocked inserts to a table within the DB2/400 database.

The following table gives a brief description of the SQL statements issued and the row descriptions. Note that the questions marks ("?") within the statements are parameter markers or variables that the client application supplies to the ODBC APIs.

Table 10.5. CA/400 Windows 3.1 Query Upload Tests						
Row Columns Rows B						
SQL Statement	Size	Per Row	Inserted	Inserted		
INSERT INTO JMBCOLL.PERF	100	10	1,000	100,000		
VALUES (?,?,?,?,?,?,?,?,?)						
INSERT INTO JMBCOLL.PERF 350	100	10	30,450	3,045,000		
ROWS VALUES (?,?,?,?,?,?,?,?,?)						

Measurement Results:

The following table shows the overall upload rates in megabytes per hour for the above SQL statements. The single insert case sends 1000 ODBC SQLExecute commands to the AS/400 server to perform 1000 inserts while the blocked insert scenario sends 87 ODBC SQLExecute commands to the server to perform 30,450 inserts.

Table 10.6. CA/400 Windows 3.1 Query Upload Performance					
SQL Insert Row Count Rows Inserted Insert Rate (MB/hr)					
1	1,000	16			
Block of 350	30,450	215			

Conclusions/Recommendations:

1. Use blocked insert when possible

Client applications that perform inserts, updates or deletes will generally perform these SQL commands one at a time to the CA/400 data access server. However, for inserts, there is an opportunity to use the blocked INSERT SQL statement which can be used to send a set of rows to the server in a single communications flow. Measurements have demonstrated that this form of insert can be over 20X faster than doing inserts one at a time.

2. Use faster clients

A large portion of upload and download operations is due to the client. Increasing the speed of the client can improve throughput.

3. Use faster communications adapters

Using slower communications adapters can result in costly delays. Upgrading the communications adapters can improve throughput.

Client Access/400 for Windows 95/NT 5250 Emulator Performance

This section shows the performance of 5250 emulation using Client Access for Windows 95/NT compared to a 5250 terminal.

Client Access provides the capability to emulate 5250 terminal sessions with the flexibility to configure the keyboard and display to the user's preferences.

5250 Emulator Performance Results

Measurement Configuration:

AS/400 200-2030 - dedicated V3R2 16 MB Memory - 4-6609 DASD (8 GB) Clients: • ValuePoint client - '486-66 Mhz - 32MB memory - Windows 95 • Pentium client - Pentium-133 Mhz - 32MB memory - Windows 95 Compared to 5250 display: 5291 model 2 CA/400 for Windows 95/NT V3R1M2 16 Mbps TR-LAN

Typical OS/400 5250 workstation screens - 24x80 resolution

Measurement Results:

Table 10.7. CA/400 Windows 95/NT 5250 Emulator Performance					
	5250 Em	ulator Performance Cor	nparison		
	CA/40	0 for Windows 95/NT V3	3R1M2		
		AS/400 - V3R2 200-2030			
(Client - '486 @ 66 MHZ	versus Pentium @ 133 M	AHZ versus 5250 Displa	У	
	Win-95				
	Win-95 '486	Pentium Resp	NT Pentium	5250 Display	
Workstation	Resp Time Time Resp Time Resp Time				
Screen	(seconds)	(seconds)	(seconds)	(seconds)	
WRKACTJOB	0.86	0.79	0.72	0.63	
WRKLIB *ALL	0.80	0.70	0.69	0.59	
WRKOBJ *FILE 0.63 0.55 0.46 0.38					
Note: Average response time (seconds)					

Conclusions/Explanations:

- 1. Client Access/400 for Windows 95/NT provides good 5250 emulator performance. NT provides slightly faster performance than Windows-95.
- 2. Faster clients provide faster response time. The Pentium times provide response times closer to the 5250 display response times.

10.4 Tips for Improving C/S Performance

Following are some tips to use when writing a client/server application which will provide the best performance:

1. Choose the right client processor.

In many queries the majority of response time is due to client processing, especially when utilizing client/server tools (eg. 4GL/CASE tools). If response time is critical, choose a fast client

processor. Figure 28 on page 0 shows the potential increase in response times for queries relative to performing the queries on a '486 @ 66MHZ client:

2. Ensure that the client is optimized for performance

Client/server applications usually require a large amount of processing power for both the client and the server. The client processor speed must be appropriate for the task. Also, memory requirements for the client may be large (toolkits, communications -- routers, buffers, etc., and operating system requirements). The response time variation between a fast well-tuned client and a under-powered client can be astounding.

3. Use the fastest communications media possible and tune the communications configuration settings

Most client/server applications tend to send a large number of requests and responses between the client application and the server. To minimize the delay due to this communications traffic, use fast media such as local area networks (LANs) to attach clients to the server. If response times are critical, do not use wide-area networks (WANs) since the communications speeds are typically measured in thousands of bits per second instead of millions of bits per second. Also, be wary of bridges, routers, and gateways since they may introduce delays when communicating across networks. Instead, keep the response time-critical clients on the same network as the server.

Use a fast PC communications adapter -- especially for file transfer operations. The communications adapter can be a major factor in constrained throughput. For download operations, a slow communications adapter can reduce throughput by over 10X compared to a fast adapter. The adapter can not keep up with the server and this results in overruns. When an overrun occurs, the server must detect the error and resend the data. This can result in large delays. If a faster adapter can not be used, communications can be tuned to reduce overruns. Following are some examples:

• TCP/IP for Windows 95/NT

When using the native Windows 95 TCP/IP communications stack, a registry entry can be changed to improve problems when using slow adapters or slow PCs. This setting can reduce performance on fast adapters and should only be changed when client adapter problems exist.

"HKEY_LOCAL_MACHINE"

Use REGEDIT to add a new string value "DefaultRcvWindow". Set the value to 4096 or decrease until retries are reduced.

When using Windows NT, the value to change is: "HKEY_LOCAL_MACHINE"

Use REGEDT32 to add the new REG_DWORD value "TcpWindowSize". Set the value to 4096 or decrease until retries are reduced.

• APPC

Set the LANWDWSTP setting from the default of 0 to 2 or greater. For slow adapters, this

will reduce the time to correct data re-transmission problems.

• Netsoft APPC Router

Consider increasing the parameter MAXDATA size from the default value of 521 to the maximum size. This value is specific to each router and can be different for each router configured. The MAXDATA size must be equal to or less than the frame size opened for the network adaptor. Increasing this value can improve performance, in particular the performance of large data transfers.

To change this parameter, open the 'Netsoft Administrator' folder, select 'Set Properties' of specific AS/400 configuration. Next select 'Properties' of the link being used (for instance 802.2). Finally, select the 'Advanced' tab.

• IPX/SPX

If you are using IPX/SPX for large file transfers the default data size sent to IPX/SPX may be increased. Create the string value:

"HKEY_CURRENT_USERAccessInternal_Components(Your Env)(Your System)IPX Max Send

The default is 1400, the maximum is 65536. Setting this value above the default may cause errors in some configurations. If problems appear after changing this value, delete the registry entry.

The frame size and buffer size for your network card should be increased to optimize the network traffic for your situation. Large file transfers perform better with larger frame sizes if your network adapter and network devices support the larger sizes. Increased buffers allow the client to offload more work to the network card. These settings can usually be controlled through the Control Panel/Network/Adapter Properties, check with your network adapter manufacturer for details. Increasing these setting will increase the system resources used by your network adapter.

Consider the following tuning tips for AS/400 communication.

 Consider increasing the Maximum Transmission Unit (MTU) from the default value of 576. The AS/400 defaults to 576 when a route is added to the configuration (via CFGTCP option 3). This value ensures packets will not be dropped over this route as all TCP/IP implementations have to support at least a 576 byte Transmission Unit.

In many cases, this value is unnecessarily small. For instance, if the route will only be used on the configured Ethernet or Token Ring, and there are no intermediate hops that only support a 576 byte packet. If this is the case, change the Route Maximum Transmission Unit size to *IFC. This will change the MTU on the Route to the Interface MTU size which defaults to the Line Description Frame size. This defaults to approximately 2000 for Token Ring and 1500 for Ethernet.

• Consider increasing the TCP receive buffer size from the default size of 8192 bytes to a larger value, for example 64384 bytes (via CFGTCP option 3).

This value specifies the amount of data the remote system can send before being read by the local application. If there are many retransmissions occurring due to the overrunning of a network adapter, decreasing this value instead of increasing it, could help performance.

• Consider increasing the TCP send buffer size from the default size of 8192 bytes to a larger value, for example 64384 bytes (via CFGTCP option 3).

This value provides a limit on the number of outgoing bytes that are buffered by TCP. If there are many retransmissions occurring due to the overrunning of a network adapter, decreasing this value instead of increasing it, could help performance.

• Refer to section Chapter 5, "Communications Performance" on page 0 for AS/400 communication tuning guidelines and specifically Section 5.2, "LAN Protocols, Lines, and IOPs" on page 0.

ANYNET support allows clients to run APPC based applications over TCP/IP. ANYNET can be considerably slower than TCP/IP and consumes more CPU than TCP/IP. Client Access for Windows 95/NT allows clients to access the AS/400 directly through TCP/IP. TCP/IP provides faster response times than ANYNET.

4. Ensure that all database requests are optimized for performance

Although the AS/400 database manager does a good job of handling database requests, it is important that performance-sensitive operations be tuned for optimal performance. Examples of database tuning are: ensure that indexes are being used, simplify SQL statements, minimize redundant operations (re-preparing SQL statements, etc.), and reduce the number of communications requests/responses by blocking. Use tools such as communications trace (STRCMNTRC), DB2/400 debug (STRDBG), Explain function (PRTSQLINF), and Performance Monitor (STRPFRMON) to assist with performance tuning.

5. Ensure that the database access method is tuned for performance

Whether the application uses DRDA, DDM, Remote SQL, DAL, JDBC, or ODBC, tuning the access method can result in performance improvement. Client Access/400 ODBC support allows each client to customize the ODBC interface to the AS/400. This is done using the ODBC.INI file for Windows 3.1 or the ODBC Administrator for Windows 95/NT. For ODBC tuning guidelines, see section "ODBC Performance Settings" on page 144.

6. If a CASE/4GL toolkit is used, tune the application for performance

Client toolkits can provide large improvements in application development productivity. But, since most are developed to communicate with multiple servers, they may not be optimized for any specific server. For more information on CASE/4GL toolkits, see section"Client/Server 4GL and Middleware".

7. Use parameter markers support when performing repetitive transactions

A parameter marker is a question mark (?) that appears in the SQL statement where a host variable could appear if the statement API string was a static SQL statement. Parameter markers enhance performance by allowing a user to prepare a statement once and then execute it many times using a different set of values for the parameter markers.

8. Reuse prepared statements

Prepares of SQL statements can take asignificant amount of time. There are two ways to reuse prepared statements:

• Only prepare statements once (using parameter markers) and use SQLExecute ODBC API

Reducing redundant prepare statements and using parameter markers instead of literals are two of the best ways to improve database server performance -- especially OLTP operations which are frequently repeated. Response time of a complex, repetitive transaction can be reduced by over 5X by changing the client application to take advantage of these improvements.

• Use package support

Package support, available with CA/400 ODBC, provides built-in reuse of prepared statements. See "ODBC Performance Settings" on page 144 for more information on configuring for package support.

9. Use stored procedures and triggers to reduce communication flows

To reduce network traffic between the client and the server and reduce response time, use stored procedures and/or triggers. Typical database serving applications send or receive from a dozen to a hundred requests/responses. Stored Procedures and triggers can reduce the number of flows significantly. Also, more processing is done at the server so the application can be completed more efficiently.

10. When possible, use SQLExecDirect for one-time execution (one flow, not two)

SQLExecDirect can replace the pair: SQLPrepare and SQLExecute. **However**, if you are doing multiple executions of the SQLStatement (looping), you should separate the SQLPrepare and SQLExecute such that the SQLPrepare is done only once and the SQLExecute is processed multiple times. This reduces both AS/400 and client processing time because the PREPARE/DESCRIBE steps do not need to be repeated. This is much more efficient than SQLExecDirect.

11. Ensure that each statement has a unique statement handle

Sharing statement handles for multiple sequential SQL statements causes DB2/400 to do FULL OPEN operations since the database cursor can not be re-used. By ensuring that an SQLAllocStmt is done before any SQLPrepare or SQLExecDirect commands, database processing can be optimized. This is especially important when a set of SQL statements are being executed in a loop. Allowing each SQL statement to have its own handle reduces the DB2/400 overhead.

12. Utilize blocking

- Use "FOR FETCH ONLY" and avoid "UPDATE WHERE CURRENT OF"
- Set maximum frame size > 2K for large upload or download

For the Windows 3.1 client, use the Global Options settings in Configuration to set the maximum frame size. For the Extended DOS client, use the TRMF setting in CONFIG.PCS.

• Use blocked inserts

Blocked Insert allows a client application to send a set of rows to the server (instead of one at a time). Measurements show that the performance of Blocked Insert can exceed 10X improvement over single row insert (eg. 1000 100-byte rows inserted)

13. Use lowest level of commitment control required

More server processing is required to process more stringent commitment control settings.

14. Define client column parameter marker variables identical to host column descriptions to allow for direct mapping on the server.

This reduces the overhead of variable type mapping.

15. Consider tuning some CASE/4GL applications (changing ODBC APIs)

Customizing "open" client applications by using the tips listed above, you may be able to improve overall performance.

16. Choose a server access method which provides high performance database serving

If your 4GL supports multiple access methods to the AS/400 server, consider the following:

a. Use ODBC for best SQL access performance

ODBC can improve performance over other SQL access methods. ODBC is the strategic database serving interface to AS/400.

b. DRDA

Distributed Relational Database Architecture (DRDA) provides acceptable performance in most cases. When possible, use static SQL statements for the best performance.

c. DDM

Distributed Data Management (DDM) does not have the flexibility of SQL but, in most cases, provides good record-level file access performance.

d. JDBC

Java Toolbox provides good C/S performance for client Java Aplications

- 17. Use client tools to assist in tuning the client application and middleware. Tools such as ODBCSpy and ODBC Trace (available through the ODBC Driver Manager) are useful in understanding what ODBC calls are being made and what activity is taking place as a result. Client application profilers may also be useful in tuning client applications and are often available with application development toolkits.
- 18. When possible, avoid extra communications layers such as AnyNet for the best performance of OLTP and large record upload/download workloads. Functions that do not require fast response times through the communications layers (e.g., ad-hoc queries and stored procedures) are a better fit for Anynet.

ODBC Performance Settings

You may be able to further improve the performance of your ODBC application by editing the ODBC.INI file on Windows 3.1. The settings in the ODBC.INI file are stored in the registry for Windows 95/NT. The recommended way to access these settings is through the ODBC Administrator in the Control Panel. The settings can be found in the registry under the Key "HKEY_CURRENT_USER.INI". The ODBC.INI file for Windows 3.1 clients contains information relating to the various ODBC drivers and data sources and is located in the Windows subdirectory for each CA/400 ODBC client. Listed below are some of the parameters that you can set to better tune the performance of the Client Access/400 ODBC Driver. The ODBC.INI performance parameters that we will be discussing are:

- Prefetch
- ExtendedDynamic
- RecordBlocking
- BlockSizeKB
- LazyClose
- LibraryView

Prefetch = choices 0, 1): The Prefetch option is a performance enhancement to allow some or all of the rows of a particular ODBC query to be fetched at PREPARE time. This option is set OFF by default. We recommend that this setting be turned ON. However, if the client application uses EXTENDED FETCH (SQLExtendedFetch) this option should be turned OFF.

ExtendedDynamic = (*choices 0,1*): Extended dynamic support provides a means to "cache" dynamic SQL statements on the AS/400 server. With extended dynamic, information about the SQL statement is saved away in an SQL package object on the AS/400 server the first time the statement is run. On subsequent uses of the statement, CA/400 ODBC recognizes that the statement has been run before and can skip a significant part of the processing by using the information saved in the SQL package. Statements which are cached include SELECT, positioned UPDATE and DELETE, INSERT with subselect, DECLARE PROCEDURE, and all other statements which contain parameter markers.

All extended dynamic support is application based. This means that each application can have its own configuration for extended dynamic support. Extended dynamic support as a whole is controlled through the use of the ExtendedDynamic keyword. If the value for this keyword is 0, no packages are used and no additional information will be added to the ODBC.INI file. If the value is set to 1 (default), when an application is run for the first time, the ODBC driver will add a line to the ODBC.INI file for the datasource in use that looks like this:

Package<Appname> = lib/packagename,usage,pkg full option,pkg not used option

Once this entry is added to the ODBC.INI file it can be modified to provide the support that the user wants.

Packages may be shared by several clients to reduce the number of packages on the AS/400 server. For the clients to share the same package, the default libraries of the clients must be the same and the clients must be running the same application. Extended dynamic support will be deactivated if two clients try to use the same package but have different default libraries. In order to reactivate extended dynamic support, the package should be deleted from the AS/400 and the clients given different libraries to store the package in. The location of the package is stored in the ODBC.INI file for Windows 3.1 and in the registry for Windows 95/NT.

Usage (choices 0,1,2): The default and preferred performance setting (2) enables the ODBC driver to use the package specified and adds statements to the package as they are run. If the package does not exist when a statement is being added, the package is created on the server.

Considerations for using package support: It is recommended that if an application has a fixed number of SQL statements in it, a single package be used by all users. An administrator should create the package and run the application to add the statements from the application to the package. Once that is done, configure all users of the package to not add any further statements but to just use the package. Note that for a package to be shared by multiple users each user must have the same default library listed in their ODBC library list. This is set by using the ODBC Administrator or by changing the ODBC.INI file.

Multiple users can add to or use a given package at the same time. Keep in mind that as a statement is added to the package, the package is locked. This could cause contention between users and reduce the benefits of using the extended dynamic support.

If the application being used has statements that are generated by the user and are ad hoc in nature, then it is recommended that each user have his own package. Each user can then be configured to add statements to their private package. For each user to have a private package, the ODBC.INI file must be modified so that each user has a different package name. Either the library name or all but the last 3 characters of the package name can be changed.

RecordBlocking = (choices 0,1,2): The RecordBlocking switch allows users to control the conditions under which the driver will retrieve multiple rows (block data) from the AS/400. The default and preferred performance setting (2) will enable blocking for everything except SELECT statements containing an explicit "FOR UPDATE OF" clause.

BlockSizeKB = (choices 1 thru 512): The BlockSizeKB parameter allows users to control the number of rows fetched from the AS/400 per communications flow (send/receive pair). This value represents the client buffer size in kilobytes (1kb=1024) and is divided by the size of one row of data to determine the number of rows to fetch from the AS/400 server in one request. The primary use of this parameter is to speed up

queries that send a lot of data to the client. The default value 32 will perform very well for most queries. If you have the memory available on the client, setting a higher value may improve some queries.

LazyClose = (*choices 0,1*): The LazyClose switch allows users to control the way SQLClose commands are handled by the Client Access/400 ODBC Driver. The default and preferred performance setting (1) enables Lazy Close. Enabling LazyClose will delay sending an SQLClose command to the AS/400 until the next ODBC request is sent. If Lazy Close is disabled, a SQLClose command will cause an immediate explicit flow to the AS/400 to perform the close. This option is used to reduce flows to the AS/400, and is purely a performance enhancing option.

LibraryView = (*choices 0,1*): The LibraryView switch allows users to control the way Client Access/400 ODBC Driver deals with certain catalog requests that ask for all of the tables on the system. The default and preferred performance setting (0) will cause catalog requests to use only the libraries specified in the default library list when going after library information.

Setting the LibraryView value to 1 will cause all libraries on the system to be used for catalog requests and may cause significant degradation in response times due to the potential volume of libraries to process.

AS/400 Memory Requirements

Multiple clients running the CPW workload were used to help determine the optimal amount of AS/400 memory needed per client. For V4R2, It was found that in the range of 2.5 to 2.8 MB per user, client response times "leveled off" such that more memory did not significantly improve response times. However, as we continued to add memory to the pool beyond 2.8 MB per user, the paging and faulting for that pool continue to decrease significantly until about 3.2 MB per user was available for each client. See "Client/Server Online Transaction Processing (OLTP)" for more information on the CPW workload.

The effects of memory depends on the kind of workload done by the jobs in the shared pool. Your memory requirements may vary depending on many factors such as high communications I/O, accessing very large database tables, frequent DASD I/O accesses, and high level of multi-processing (sharing) in the pool.

Note that the AS/400 Work Management manual (SC41-3306) has a good set of recommendations on memory tuning. Knowledge of these tips can assist in building high-performance system environments.

Chapter 11. Domino for the AS/400

Performance information for Lotus Domino for AS/400 is included in this section. Domino for AS/400 provides a variety of functions, with this section limited to the performance of the mail server function.

There are many factors that can impact overall performance (e.g., end-user response time, throughput) in the AS/400 Domino environment, some of which are listed below:

- AS/400 processor speed
- utilization of key AS/400 resources (CPU, IOP, memory, disk)
- object contention (e.g. mutex waits, lock waits)
- speed of the communications links
- congestion of network resources
- processing speed of the client system

The primary focus of this section will be to discuss the performance characteristics of the AS/400 as a server in a Domino environment, providing capacity planning information and recommendations for best performance.

11.1 Workload Description:

• Mail

The mail workload scenario was driven by an automated environment which executed a script similar to the mail workload from Lotus NotesBench. Lotus NotesBench is a collection of benchmarks, or workloads, for evaluating the performance of Notes servers. The results shown here are not official NotesBench measurements or results. The numbers discussed here may not be used officially or publicly to compare to NotesBench results published for other Notes server environments. For official AS/400 audited NotesBench results, see http://www.notesbench.org. (Note: in order to access the NotesBench results you will need to apply for a userid/password through the Notesbench organization. Click on Site Registration at the above address.)

Each user completes the following actions an average of every 15 minutes:

- ✤ Open mail database
- ✤ Open the current view
- ✤ Open 5 documents in the mail file
- ✤ Categorize 2 of the documents
- Compose 2 new mail memos/replies
- ✤ Mark several documents for deletion
- ✤ Delete documents marked for deletion
- ✤ Close the view

11.2 Domino for AS/400, Release 5.0.

Release 5.0 of Domino (R5 for short), included a major initiative to improve Domino performance. Some of the changes include:

- 1. Redesign of the on disk structure (ODS) for storing databases.
- 2. Memory and I/O optimization
- 3. Ability to use multiple mail boxes.
- 4. Transaction logging to improve server recovery time and data integrity.
- 5. The ability to assign individual priorities to tasks within a server (see section 11.7).

These changes and others helped reduce the number of disk I/Os, and in certain instances, the amount of CPU time used. The biggest factor in determining the benefit of R5 over 4.6 is the number of active users in a partition. For a large number of active simple mail users in a partition (above 1800), R5 will begin to see performance improvements over 4.6. At lower numbers of users in a partition, R5 does not show any improvement in the tests we ran, and in some cases, R5 showed a slight degradation in performance due to an increase in CPU time.

However, even below 1800 users in a partition, R5 still showed performance improvements at critical times. For example, during rampup (users are connecting to the Domino for AS/400 server), R5 showed a greatly improved response time. This improvement in connection time held true even up to 7000 users in a single partition. In 4.6, it was difficult in our test environment to get beyond 2600 users.

The net result of all this is that if you have many partitions with low number of users, you now have the option of combining those into a single partition (up to 7000 active simple mail users), if you have the available CPU resource. This is a tradeoff between manageability and performance, because as you add users to a partition, the CPU cost per user increases (from 1200 to 7000 the cost goes up 40%).

The following graph shows how the processor costs will increase as the number of users in a partition increases. This increase is due to management overhead and contention for Domino resources. Notice that 4.6 and R5 have about the same cost at around 1800 users in a partition. However, the cost for 4.6 increases rapidly as users are added to the partition. R5 does not show as large of an increase.



Figure 11.2 CPU cost per Domino Mail User.

Additional conclusions that we can draw from the R5 data are as follows:

- 1. We ran our tests with the NSF_Buffer_Pool_Size_MB set to 300. This reduced the CPU utilization only a small percentage, but reduced the page faulting demands significantly when compared to measurements set to 507MB. This certainly seems counter-intuitive. The memory management changes in R5 have changed how the NSF buffer pool is used, so the behavior we experienced is not fully understood. You may need to experiment with your settings to determine an optimum value.
- 2. Transaction logging adds from 5% to 7% CPU cost for the environment that was measured (1800 total simple mail users in 2 partitions). The total number of disk I/Os were reduced slightly, but the overall disk utilization went up slightly (related to item #2 below). The conclusion to be drawn is that the CPU and disk costs are justified if server reliability and recovery speed are considered.
- 3. While R5 caused an increase in disk utilization over 4.6, the actual number of disk reads and writes were reduced by 5%. The disk utilization increased because the disk write cache was being overwritten due to the increase in the size of the disk writes(10K vs 7K in our Simple mail workload) and the way the disk writes were not spread evenly over time. V4R4 will compensate for most of this with its IFS synch daemon redesign. For V4R3, you will need to ensure you have enough disk configured to keep the disk % busy below guidelines.
- 4. Running the simple mail user tests with R5 clients vs. 4.6 clients showed no performance difference.
- 5. When we ran the 3650 user tests in 2 partitions, the clients in the R5 test were able to sign on 2.5 times as fast as the 4.6 users. This is due to the resource management improvements in R5.
- 6. Specific tests to analyze the improvement of using multiple mail.box'es were inconclusive. However, these tests were run with 900 users in a partition, which may not be high enough to show the benefits of multiple mail.boxes. The 7000 user test used 4 mail.boxes, and ran very well. A 7000 user run with a single mail.box was not completed. However, knowledge mail routing algorithms combined with the tests we ran suggest that 1 mail.box is sufficient for low number of users in a partition, 2 is sufficient for most environments, and more than 2 does not improve performance significantly, but still may be beneficial, up to 4 mail.boxes.
- 7. While V4R4 for 4.6 shows a 10% CPU improvement over V4R3, the benefit is not as great in R5 because R5 reduced some of the same bottlenecks that V4R4 reduced. Therefore, for R5, V4R4 has about a 5% improvement over V4R3.

11.3 V4R4 changes that affected Domino Mail Serving Performance:

- 1. The OS/400 Integrated File System improved the way it writes changed files to disk. Instead of waking up every so often and forcing changed files to disk, the files are aged appropriately and written out only when needed in order to ensure the files have a minimal exposure time. The Integrated File system is also more tightly integrated with OS/400 storage management to take advantage of the storage management aging process. The result is a 40% reduction in I/O and a corresponding CPU reduction for the simple mail workload. This environment was dominated by write operations (99%). Domino environments with a lower percentage of write operations may not realize as much benefit from the Integrated File System changes described above.
- 2. Domino makes heavy use of timer events, and was stressing the OS/400 timer algorithms more than any other application had done before. For this reason, the OS/400 timer management code

was optimized to handle this heavy volume of timer events. Domino R5 also improved their use of timer events, which is why V4R4 shows less of an improvement over V4R3 than Domino 4.6.

3. The net result of V4R4 is a 10% reduction in CPU costs and a 40% reduction in I/O for the simple mail workload for Domino 4.6. For Domino R5, the CPU improvement is about 5% while the disk improvement is still 40%.

11.4 AS/400e Dedicated Server for Domino:

Available 9/24/99, the Dedicated Server for Domino (DSD) models deliver exceptional price performance for "Domino-only" AS/400 environments! The AS/400e Server 170 will have 3 new features: 2407, 2408, and 2409. The DSD models are intended for Domino-only environments such as:

- e-mail
- calendaring and scheduling
- web serving
- standard Lotus Domino template applications (discussion database, workflow, etc.) and custom developed applications written with Domino designer which perform no external calls, relational database access, or Java integration

The Domino-only environment which exhibits the best performance can be compromised if :

- client/server processing exceeds 10-15% of the CPU capacity
- interactive processing exceeds reasonable system administration activities
- application integration functions (e.g. DB2 Universal Database access, external program calls, and Java) comprise more than approximately 25% of the work being done

DSD models running Domino applications will accommodate a small amount of the use of application integration functions without any loss to efficiency of the Domino-only environment. Should the integration functions exceed approximately 25% of the workload, the Domino-only environment will be compromised and performance will begin to decrease. Monitoring the CFINT CPU usage will be the easiest way to determine whether a given application environment is suitable for a DSD model. If large amounts of CFINT CPU processing occur, it indicates an excessive amount of interactive and/or application integration function is occurring and degrading the efficiency of the system. For additional information on the behavior of the DSD models with respect to the Interactive and Processor CPW limits, please refer to Chapter 2, "AS/400 RISC Server Model Performance Behavior".

From the data in Table 11.1, note that the more economical DSD model 170-2409 when maintaining a Domino-only environment can deliver similar mail serving capacity as a model 170-2388, both running V4R4 and Domino R5. However, for work other than Domino-only, the model 170-2409 will have much less capacity than the model 170-2388.

Please refer to <u>http://www.as400.ibm.com/lotus_notes/domsupport.htm</u> for details on PTFs and QMU levels required for DSD models.

11.5 Mail Serving Performance Conclusions/Recommendations:

- 1. Domino R5, 4.6 and 4.6.2 running on OS/400 V4R2+ provides an industry leading scaleable mail server. The AS/400's unique subsystem architecture allows multiple Domino partitions on a single piece of hardware. Not only does this improve performance, but it allows for better manageability of groups of users.
- 2. The 170-2292 mirrored system with 1GB of main storage had very good response time at only 77% CPU utilization. This system had 1 drive in ASP 1 for OS/400, and 2 mirrored drives in ASP 2 for Notes data (mailboxes, name and address book, etc.)
- 3. The 170-2292 RAID-5 system with 1GB of main storage also had very good response time but could only support 1300 users with the 4 drives protected with RAID-5. The 3 additional I/Os for each write due to RAID-5 protection limited the number of users this configuration could support. Configuring additional arms would have relieved this bottleneck.
- 4. The S40-2208 with 40GB of main storage demonstrates how easily Domino for AS/400 scales upwards to large numbers of users and many partitions. Even at 27,000+ users, the system still had available resources for more Domino users or other applications.
- 5. The S40-2261 with 20GB of main storage ran at only 57% CPU utilization, and paging was very low. This means valuable main storage and CPU are available for other users and/or mission critical applications.
- 6. The S40-2261 main storage was reduced to 12GB with minimal impact on response time, CPU utilization, or disk utilization. The faulting rate was still very acceptable.
- 7. The 150-2269 also has leftover CPU and main storage capacity when running 30 active users. Even the single disk arm on this system was not overly utilized at 25%.

11.6 Domino Performance Tips/Techniques:

- 1. As the number of active users in a partition increases, contention for resources will cause an increase in CPU consumption for each user (see figure 11.2). Best results were achieved when the number of active users was limited to 1000, even though the S40-2208 with R5 successfully ran 7000 users in a partition, and the 170-2292 with 4.6 successfully ran 1350 users in a single partition.
- 2. Initial user connection places the heaviest load on the AS/400, requiring the largest amount of CPU, main storage, and disk resources. When sizing a system, ensure there is sufficient capacity for this activity as well as the typical peaks of activity throughout the day. R5 reduces the contention during initial signon, but does not eliminate it.

3. For the system storage pool in which the Domino server and users run, you will need to configure at minimum 7MB + 600KB per active user for small systems and 8MB + 800KB per active user for large systems. Note this does not include storage required for other system storage pools (e.g. *MACHINE, *SPOOL, etc.). Adding more main storage than what is recommended here will provide even better response times and will also provide capacity for future growth or workload peaks.

Follow the faulting threshold guidelines suggested in the Work Management guide by observing/adjusting the memory in both the machine pool and the pool that the mail servers run in.

- 4. AS/400 notes.ini / server document settings:
 - For the first server job created, additional threads will be created as needed until 100 is reached. Then the secondary server job will be created with 100 threads all at one time. If this number for secondary threads is too high, you may either notice a general slowdown on the system as the entire set of threads is created, or you may get server errors if there are more than 2048 files open in a single job. Depending on your configuration, these errors begin to show up if the number of threads per job approaches or exceeds 400. To set the initial and secondary threads to 100, put the following lines in the server's NOTES.INI file.
 - ✤ SERVER_MAXINITIALTHREADS=100
 - ✤ SERVER_SECONDARYTHREADS=100
 - Mail.box setting.

Setting the number of mail boxes to more than 1 may reduce contention and reduce the CPU utilization. Setting this to 2, 3, or 4 should be sufficient for most environments.

MAILMAXTHREADS

This is the maximum number of threads that the mail router job can create. The default is one thread per server port. Increase this number to improve mail- routing performance, which is especially important for mail hub servers. Use the Show Server command from the Domino Console to check for pending mail and increase MAILMAXTHREADS by 1 until the pending mail typically shows 0 or reaches an acceptable level for your environment.

• NSF_Buffer_Pool_Size

This controls the size of the memory section used for buffering I/Os to and from disk storage. The best results for the 1350 user measurements were obtained when this was set to 507000000. If you make this too small and more storage is needed, Domino will begin using it's own memory management code that adds unneeded overhead since OS/400 already is managing the virtual storage. The 507000000 value allows for the maximum amount of memory for the NSF Buffer Pool and the memory is not allocated until needed. Use the Show Server command from the Domino Console to observe the maximum value a server has used for the NSF Buffer Pool. Do not set this above 508000000, since doing so will cause Domino to calculate its own value, which may not be optimal.

For R5, a value of 300 for NSF_Buffer_Pool_Size_MB showed a slight improvement over a size of 507. Also, R5 allows settings above 508, even though for our configuration this did not make any performance difference.

- 5. AS/400 environment variable settings.
 - Notes_SHARED_DPOOLSIZE. Set to 12000000, or don't set it because this is the AS/400 default value. This value controls how Domino memory management is done. By increasing this value, you reduce the amount of management overhead. Do not make this variable larger than 12000000.
 - Notes_AS400_CONSOLE_ENTRIES set to 10,000 (the default). This is the size of the console file that displays the status messages when you enter the DSPDOMCSL or WRKDOMCSL commands. As this file grows, the response time for these two commands increases.

For more detail on the above settings, see the Domino Server Administrator Guide.

6. Use *LOOPBACK to prevent data from going out over network.

For servers communicating on the same system, you can shortcut the communications path by using *LOOPBACK instead of transmitting the data out of the system just to have it come back again. See <u>http://www.as400.ibm.com/techstudio</u> for details.

7. Dedicate servers to a specific task

This allows you to separate out groups of users. For example, you may want your mail delivered at a different priority than you want database accesses. This will reduce the contention between different types of users. Separate servers for different tasks are also recommended for high availability.

8. MIME format.

For users accessing mail from both the Internet and Notes, store the messages in both Notes and MIME format. This offers the best performance during mail retrieve because a format conversion is not necessary. NOTE: This will take up extra disk space, so there is a trade-off of increased performance over disk space.

9. Full text indexes

Consider whether to allow users to create full text indexes for their mail files, and the use of them whenever possible. These indexes are expensive to maintain since they take up CPU processing time and disk space.

10. Replication.

To improve replication performance, you may need to do the following:

- Use selective replication
- Replicate more often so there are fewer updates per replication
- Schedule replications at off-peak hours
- Set up replication groups based on replication priority. Set the replication priority to high, medium, or low to replicate databases of different priorities at different times.
- 11. Additional references

The following web site contains additional Domino information and white paper resources. See <u>http://www.ibm.com/as400/developer/domino/</u> then click on performance. Some of the information will be redundant to what is provided in this document:

11.7 Domino Subsystem Tuning

The objects needed for making subsystem changes to Domino are located in library QUSRNOTES and have the same name as the subsystem that the Domino servers run in. The objects you can change are:

- Class (timeslice, priority, etc.)
- Subsystem description (pool configuration)
- Job queue (max active)
- Job description

The system supplied defaults for these objects should enable Domino to run with optimal performance. However, if you want to ensure a specific server has better response time than another server, you could configure that server in its own partition and change the priority for that subsystem (change the class), and could also run that server in its own private pool (change the subsystem description).

New for R5, you can create a class for each task in a Domino server. You would do this if, for example, you wanted mail serving (SERVER task) to run at a higher priority than mail routing (ROUTER task). To enable this level of priority setting, you need to do two steps:

- 1. Create the classes that you want your Domino tasks to use.
- 2. Modify the following IFS file '/QIBM/USERDATA/LOTUS/NOTES/DOMINO_CLASSES'. In that file, you can associate a class with a task within a given server.
- 3. Refer to the release notes in READAS4.NSF for details.

11.8 Mail Serving Capacity Planning

From the measurements listed above, we can determine that each CPW available on an AS/400 processor can support approximately 5.4 light mail users at a 70% CPU utilization. So to determine the 70% capacity for light mail for any AS/400 processor, you can multiply its CPW by 5.4. This is a rough estimate since mail serving performance does not scale exactly to CPWs. Also, when you size a system, a typical mail user is usually about 3 times as complex as the light mail users measured above. Therefore, if you need to do a detailed capacity planning exercise for Domino on AS/400 please refer to the following sources of information:

- The Workload Estimator will estimate the proper sized AS/400 for Domino, Java, Net.Commerce and traditional workloads, individually or in combination. See **Appendix B**, AS/400 Sizing
- <u>Http://www.as400.ibm.com/lotus_notes/notes.htm</u> under the section 'Domino Sizing Information.' This information will be removed in the near future, so you should use the Workload Estimator above if possible.

11.9 Mail Serving Performance Measurements:

The following tables provide a summary of the measured performance data. These charts should be used in conjunction with the rest of the information in this section for correct interpretation. Results listed here do not represent any particular customer environment. Actual performance may vary significantly from what is provided here.

 Table 11.1. Simple Mail Serving Performance Data

 Mail Serving With Domino on AS/400 Server Models

Madal	# Active Notes	# Domino	Main	Response Time	CPU	# Disk	Disk
Widder	Users	raruuons	Storage	(secs)	70 Dusy	Arms	- 70 Busv
Domino 5.01							2005
DSD Models							
170-2407	1,250	1	1.0GB	0.1	70.1	10	6.4
V4R4							
RAID-5							
170-2408	2,000	2	3.5GB	0.055	65	10	7.3
V4R4							
KAID-5	2 650	2	2.5CP	0.051	62.1	10	22.0
170-2409 V/R/	3,030	5	5.50B	0.031	03.1	10	22.9
RAID-5							
Domino 5.0							
170-2388	3,650	3	3.5GB	0.218	65.6	10	32
V4R3	r -						
RAID-5							
170-2388	3,650	3	3.5GB	0.082	63.1	10	18.7
V4R4							
RAID-5							
170-2388	1,800	2	3.5GB	.040	32	10	10.4
V4R3							
RAID-5							
Iransaction							
Logging	1 200	2	2.5CD	0.028	20	10	10
170-2500 VAP3	1,800	2	3.30b	0.038	50	10	10
RAID-5							
S40-2208	7 000	1	40Gb	0.2	29.3	228	0.4
V4R3	,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,	-	1000	0.2			0.1.
RAID-5							
Domino 4.6							
170-2388	3,650	3	3.5GB	0.084	58.9	10	15.2
V4R3							
RAID-5							
170-2388	3,650	2	3.5GB	0.08	67.4	10	20
V4R3							
RAID-5	10.290	20	4000	NI/A	42	116	1.4
740-2070 VADA	19,380	50	4000	IN/A	45	110	1.4
RAID-5							
740-2070	19 380	30	40GB	N/A	48	170	19
V4R3	19,500	50	4000	1.071	40	170	1.9
RAID-5							
S40-2208	27,030	30	40GB	.082	77	116	3.8
V4R3							
RAID-5							
170-2292	1,350	1	1GB	.270	77	3	36
V4R2							
Mirrored							
170-2292	1,300	1	1GB	.172	71	4	NA
V4R2							
RAID-5							
S40-2207	17,600	16	40GB	.126	74	96	3
V4R2							

Model	# Active Notes Users	# Domino Partitions	Main Storage	Response Time (secs)	CPU % Busy	# Disk Arms	Disk % Busy
S40-2261 V4R2	10,400	12	20GB	.050	57	47	12
S40-2261 V4R2	10,400	12	12GB	.089	57	47	13
150-2269 V4R2	30	1	64MB	.349	25	1	15
Note [.]							

Data shown above should not be compared to audited Notesbench results.Results may differ significantly from those listed here.

• These measurements are not meant to be interpreted as maximum user data points.

Chapter 12. OS/400 Integration of Lotus Notes Performance

This section includes measurement data using the 6616 Integrated PC Server - on both the V3R2 CISC platform as well as the V3R7 Enhanced RISC platform. To use the AS/400 CPU data that is provided in this section to estimate utilization for other AS/400 models, please refer to Chapter 2, "AS/400 System Capacities and CPW" on page 0 for information on relative performance of AS/400 models. Due to differences in IPCS support, one should use CISC measurement values to estimate/project to other CISC models. Likewise, RISC values should be used to estimate other RISC models.

The V4R1 PRPQ (P84304) - IPCS EFS (Enhanced File System) further improves the performance capability of the IPCS - particularly those with more than 64MB of memory. It is able to manage/use large amounts of IPCS memory(RAM) to provide efficient file/data caching service to OS/2 and applications such as Notes. EFS is available for all models of the IPCS including those used in V3R2, V3R7, and V4R1 configurations. (**Note** A primary component of EFS is a version of OS/2 HPFS386.)

Performance information for OS/400 Integration of Lotus Notes(**) is included in this section. The information will be divided into two general areas: number of Lotus Notes clients that can be supported, and performance guidelines for DB2 Integration function.

For a complete overview and understanding of OS/400 Integration of Lotus Notes, please refer to the following resources:

- OS/400 Integration of Lotus Notes (SC41-3431-02) V3R7Enh
- OS/400 Integration of Lotus Notes (SC41-3431-01) V3R2
- Using Lotus Notes on the Integrated PC Server (IPCS) for AS/400 (SG24-4779-00). This is a Redbook publication.

The File Serving IOP (FSIOP) has been renamed the Integrated PC Server (IPCS). These terms are used interchangeably in this section and refer to the following models:

- IPCS Type 6506 (based on i486 66mHz CPU; RAM size up to 64MB)
 - 6506 was previously released and is compatible with AS/400 RISC and CISC models (except Model 150)
- IPCS Type 2850 (based on Pentium 133mHz CPU; RAM size up to 128MB)
 2850 is available only on AS/400 Models 150
- IPCS Type 6616 (based on Pentium 166mHz CPU; RAM size up to 256MB)
 6616 is compatible with AS/400 RISC and CISC models (except Model 150)

OVERVIEW: Substantial improvements are seen on the IPCS type 6616 as compared to the 6506:

- 480 Mail users versus 200 (non-EFS)
- 700 Mail users (6616 with EFS)
- DB2 Integration Import rates (non-EFS)**
 - ✤ 2MB/minute vs .5MB/minute (2.1KB recordlength)
 - ✤ 2.8MB/minute vs 1MB/minute (6.2KB recordlength)
- DB2 Integration Shadowing rates (Inserts, Updates, Deletes) (non-EFS)**
 - 100 changes/minute vs ~15 changes/minute (a 10,000 document Notes database)
 - ✤ 62 changes/minute vs ~7.5 changes/minute (a 20,000 document Notes database)

**Note: Measurement data is not available for these DB2 Integration benchmarks with EFS. No performance improvements are expected from EFS caching since all data would be 'new'. However there may be some improvements due to more efficient I/O handling by the EFS.

12.1 Number of Notes Clients Supported

The Notes server on an IPCS can support different numbers of users depending on the type and number of requests from the Notes clients. Three workload scenarios are described below and measurement results using these scenarios are shown in the tables following the descriptions.

Workload Scenario Descriptions:

• Mail

The mail workload scenario was driven by an automated environment which executed a script similar to the mail workload from Lotus NotesBench. Lotus NotesBench is a collection of benchmarks, or workloads, for evaluating the performance of Notes servers. The results shown here are not official NotesBench measurements or results. The numbers discussed here may not be used officially or publicly to compare to NotesBench results published for other Notes server environments.

Each user completes the following actions an average of every 15 minutes:

- Open mail database
- Open the current view
- ✤ Open 5 documents in the mail file
- Categorize 2 of the documents
- Compose 2 new mail memos/replies
- ✤ Mark several documents for deletion
- Delete documents marked for deletion
- Close the view

• Mail and Discussion

The Mail and Discussion workload scenario was driven by the same (as Mail) automated environment described above. It executes a script containing representative scenarios/sequences similar to Lotus NotesBench's. Again, these numbers are not to be used officially or publicly to compare to NotesBench results published for other Notes server environments.

Each user completes the following actions an average of every 15 minutes:

- ✤ Open mail database
- ✤ Make sure 50 documents exist
- ✤ Open the current view
- ✤ Open 5 documents in the mail file
- Categorize 2 of the documents
- Compose 2 new mail memos/replies

- ✤ Mark several documents for deletion
- Delete documents marked for deletion
- Open a discussion database
- ✤ Make sure 200 documents exist
- ✤ Open the current view
- Page down the view 2 times
- Set the unread list to a randomly selected 30 documents
- Open the next 3 unread documents
- ✤ Close the view

• Mail and Discussion with Import

For the mail and discussion with import workload scenario, three import requests (importing data from DB2/400 to a Notes database) were performed during the measurement. The imports executed sequentially and an import was active during the entire measurement. Please see "Lotus Notes DB2 Integration Performance" for additional information on the performance of import function. The mail and discussion portion of this workload scenario was the same as the mail and discussion workload scenario described above.

The tables below includes the data/measurement values that were published in the prior release of this document. They provide performance information and guidelines using the three workload scenarios described above:

Table 12.1. Memory Guidelines for Mail Workload				
Memory Guidelines for Number of Notes Mail Users Supported on Integrated PC Servers				
Memory Size	Max Number of Users			
32MB	100			
48MB	150			
64MB	200			
128MB	400 (w/EFS)			
256MB	700 (w/EFS)			

IPCS Memory Guide for Mail Workload

Observations: Memory requirements by the Mail workload is proportional to the number of users regardless of the IPCS's processing speed.EFS is recommended for IPCSs with greater than 64MB. For 64MB or less, the Mail workload results indicate no performance gains.

A key resource in this environment is the internal cache for database (.NSF file) activity. This buffer pool size is specified with the NSF_BUFFER_POOL parameter in the NOTES.INI file which becomes effective at the startup of the Notes server application (STRNWSAPP). To monitor the actual use of this cache, use the Show Stat Database command at a Remote Console.

Here's a sample response from Show Stat Database:

```
COMMAND SENT: show stat database
 Database.BufferControlPool.Peak = 784872
 Database.BufferControlPool.Used = 659592
 Database.BufferPool.Maximum = 264241152
                                             <-- from notes.ini
 Database.BufferPool.Peak = 33926400
                                             <-- peak use so far
 Database.BufferPool.PerCentReadsInBuffer = 97
 Database.BufferPool.Reads = 52172
 Database.BufferPool.Used = 33760468
                                             <-- current use
 Database.BufferPool.Writes = 80848
 Database.NIFPool.Peak = 1308120
 Database.NIFPool.Used = 823933
 Database.NSFPool.Peak = 1242714
 Database.NSFPool.Used = 797506
```

The value shown for Database.BufferPool.Maximum is from default value specified in NOTES.INI or the available physical memory - whichever is smaller. We recommend that the Database.BufferPool.Used & .Peak be monitored regularly. '.Maximum', should be increased when '.Peak' approaches its value.

The Operating System for the IPCS is OS/2 Warp Connect with the HPFS dasd file system. A file system provides varying degrees of caching, concurrent open files, files sharing and file locks. HPFS in our case provides for a 2MB maximum disk cache - a rather modest size.EFS will greatly improve the HPFS dasd file system performance by expanding the cache size along with improved caching algorithms. I/O rates, a key indicator of caching effectiveness, has been reduced to about half. In the IPCS environment, the AS/400 CPU is used mostly to service dasd I/O requests. Thus when I/O rates are reduced, the AS/400 CPU utilization is reduced correspondingly. There is also significant reduction of the IPCS's CPU utilization.

Suggested Technical References

An important share of this OS/400 product is based on IBM Austin's work on Notes for the OS/2 PC platform. As such it may be instructive to study some of their technical papers regarding their experience and recommendations on the product set. These papers cover the 'Entry' (using Warp Connect with HPFS) and 'High End' (using Warp Server Advanced with HPFS386). Our IPCS product would correlate to the 'Entry'. Be aware of these differences when reading and/or applying the information provided.

IBM Austin's papers may be accessed at these URLs:

- Http://www.austin.ibm.com/pspinfo/noteonws.html#contents
- <u>Http://www.austin.ibm.com/pspinfo/noteonws1.html#flex</u>
- <u>Http://www.austin.ibm.com/pspinfo/noteonws3.html#usage</u>

At the third URL, there is background information and also recommendations on 'Tuning Notes Server Memory Usage'.

Please check the following Web pages for additional technical information, support and update notifications.

- <u>Http://www.as400.ibm.com/products/software/notes/performance.htm</u>
- <u>Http://www.as400.ibm.com/notes/performance</u>

AS/400 Machine Pool Guidelines

The following guidelines can be used to estimate the amount of AS/400 memory that will be required in the machine pool (System Pool #1) to support the indicated Integrated PC Server:

- IPCS 6616 1.4MB plus 1.8MB per port ٠
- IPCS 6506 1.4MB plus 1.8MB per port ٠
- IPCS 2850 4.5MB ٠

IPCS 6506 and 6616 on CISC

Table 12.2. Users on Integrated PC Server Running Notes						
6506 on CISC (non-EFS)						
	Lotus 1	Notes on IPCS-6506 64MB	RAM			
	Vario	ous AS/400 Models Used, V	/3R2			
AS/400 Int PC						
Workload	Number of	Model(s)	AS/400	Server (6506)		
Scenario	Users	CPU (%)	IO/Sec	CPU (%)		
Mail	100	5% 20S-2010	38	33		
Mail	200	9% 20S-2010	68	66		
Mail + Discussion	60	4% E25	18	33		
Mail + Discussion	60	13% E25	70	70		
with Import						
Note:						
Average utilization of 16	Mbps TRLAN was <1%	o for all scenarios.				

Table 12.3. Users on Integrated PC Server Running Notes					
6506 on CISC (non-EFS)					
Lotus Notes on IPCS Type 6506 w/64MB RAM					
AS/400 model E25 Used, V3R2, 13 DASD Arms, 64MB					
				Int PC	
Workload	Number of	AS/400-E25	AS/400-E25	Server (6506)	
Scenario	Users	CPU (%)	IO/Sec	CPU (%)	
Mail	100	7	32	30	
Mail	200	17	66	57	
Note:					

• Average utilization of 16Mbps TRLAN was <1% for all scenarios.

• Virtual DASD (VDASD) support for IPCS on CISC is enhanced with multiple Storage Management tasks (3 for 6506; 7 for 6616)

Table 12.4. Users on Integrated PC Server Running Notes						
6616 on CISC (non-EFS)						
	Lotus Notes	s on IPCS Type 6616 w/2	56MB RAM			
	AS/400 model	E25 Used, V3R2, 13 DAS	SD Arms, 64MB			
Int PC						
Workload	Number of	AS/400-E25	AS/400-E25	Server (6166)		
Scenario	Users	CPU (%)	IO/Sec	CPU (%)		
Mail	100	9	38	8		
Mail	200	17	67	15		
Mail	300	25	104	24		
Mail	400	34	142	36		
Note:						
• Average utilization of 16Mbps TRLAN was <1% for all scenarios.						
• Virtual DASD (VDASD) support for IPCS on CISC is enhanced with multiple Storage Management tasks						
(3 for 6506; 7 for 661	(3 for 6506; 7 for 6616)					

Observations: Tables 12.2, 12.3, and 12.4 above summarize data for 3 CISC combinations of 6506 and 6616.

- IPCS 6506 on CISC
- IPCS 6506 on CISC
- IPCS 6616 on CISC

Note that the 6616 configuration has substantially extended the maximum number of users supported from 200 to more than 400. Although not listed in in Table 59, up to 480 users have been observed with this configuration. The main factors for this improvement are:

- Pentium/166mHz vs i486/33mHz; this represents processing speed ratio of about 3.5 times. This is confirmed by the IPCS's CPU%: e.g. on E25 at 200 users 6506 runs at 57%; 6616 is at 15%.
- IPCS (with large RAM): Configurations with greater than 64MB should benefit greatly from installing EFS. At this time there is no direct measurement data (of IPCS on CISC with EFS) but it is expected that the pivotal parameters (DASD I/O rates, IPCS CPU%) will parallel the improvements shown in Table 64 on page 163 and Table 65 on page 164 (IPCS on RISC). We expect that the E25/6616 configuration of Table 59 above will support Mail users up to the 700 level with the E25 CPU utilization in the 40% range and the IPCS CPU in the 50% range.
- VDASD performance on CISC was enhanced by the PTF# NF14527, (on Cum. tape# 7063) which supported and used multiple tasks (3 for the 6506; 7 for 6616). See "Multiple VDASD Tasks for CISC" section below.

The AS/400 server model 20S-2010 shown in Table 12.2 above has a 5.9 Relative Processor Rating (RPR) for non-interactive work. Using the data for the 100 user mail workload, we can predict that this same workload would require approximately 6.6% CPU on an AS/400 model E25 which has an RPR rating of 4.0. **Note** This estimate is now validated (as correct) with data shown in Table 12.3 above. See Chapter 2, "AS/400 System Capabilities and CPW" for additional information on RPR values to enable CPU comparisons for other AS/400 models.

All of the AS/400 CPU required to support both the Notes server and DB2 Integration function runs in non-interactive mode on server models. Server models have a better price performance than traditional models for this environment.

**** Note **** that for a given workload, differences in the rates shown in the tables are due to workload and measurement variability. These differences should not affect their usefulness as input to this level of system performance analysis.

****Also **** these CISC values and rates should be used to project performance onto other CISC models. Likewise use RISC values to project for other RISC models.

Multiple VDASD Tasks for CISC

The VDASD path has been enhanced for the support of IPCS on CISC. In the original IPCS release, there was a single VDASD task to perform all the DASD requests from the IPCS. With the multiple task enhancement, each 6506 will have three(3) of these tasks initiated at VARY ON time. Similarly, each 6616 will have seven(7).

Discussion: The service time of a vdasd task is basically that of the DASD service time - let's say ~ 15 milliseconds. This implies that a single task cannot perform more than about 60 to 70 I/Os per second in support of an IPCS. Response times would be extreme due to queuing times at such a max'd out server.

The limitation of the single VDASD task (original release) was masked by the 6506's (relatively low) speed and memory capacity. With the 6616 whose speed and memory capacity is about tripled - a higher VDASD bandwidth was required. The multi-task VDASD support will perform well for the 6616 and for some workload environments, may also improve the performance characteristics of the 6506.

Performance Data: The AS/400 Performance Tools can be used to create reports from data collected by the AS/400 system performance monitor (STRPFRMON). The Component report shows IPCS utilizations in the "IOP Utilizations" section, and shows job and task related data in the "Job Workload Activity" section.

The following section discusses the IPCS supporting tasks/names that will be reported and listed in the Performance Monitor reports.

Tasks related to IPCS processing

For CISC, the following names/types of tasks are used to process the DASD I/O requests for an IPCS:

- ROUTxx handles communications across bus to IPCS
- #O00yy DASD I/O server task to map HPFS space to AS/400 DASD space
- SM0nyy storage management task to perform physical I/O to DASD

There will one of each of the first two task types, and multiples of the latter - where 'n' will be 1, 2, and 3 for the 6506; and 1, 2, ... 7 for the 6616. There are additional tasks with similar names which are used for IPL and administrative processing.

For RISC, the following names/types of tasks are used to process the DASD I/O requests for an IPCS:

- FPHA-NWSDname performs function similar to ROUTxx and #O00yy above
- SMDSTASKaa multiple tasks (2 x # of DASD arms; with a maximum of 32 tasks) are used to process DASD I/O for all IPCSs on the system.

There will be a FPHA-NWSDname task for each IPCS on the system, and the SMDSTASKaa tasks are shared by all IPCSs on the system. Additional tasks, FPHI-NWSDname and FPN-NWSDname, also exist and are used for IPL and administrative processing.

For the scenario involving import, additional AS/400 jobs are used to retrieve data from DB2/400 for import to the Notes database. Jobs used for this function will have names similar to QZDASOINIT and QZRCSRVS.

Mail and Discussion with Import workload scenario

This scenario performed significantly more IO writes per second than the mail and discussion without import scenario. The high number of writes in this scenario occurred as a result of importing the data from DB2/400 and creating new documents in the Notes database. Having an import active used significantly more AS/400 and IPCS resources and impacted response times for the 60 attached users. See 12.2, "Lotus Notes DB2 Integration Performance" on page 164 for additional information and recommendations regarding the import function.

Multiple IPCS on CISC (non-EFS)

The following tabulation was provided in V3R2. Measurement data of this configuration/environment is not available for the 6616 IPCS. However, if required, one could project by scaling based on the preceding tables. For example if we were to replace the 6506s with 6616s, we would expect:

- each 6616 could support at least 400 users; total of 2400;
- each 6616's CPU would be at about 36%
- I/Os would rise to 800+ per second; assuming sufficient DASD arms;
- CISC CPU would be at about 15%.
- Previous additional measurements with eight(8) 6506s achieved expected scalability.

Table 12.5. Multip	Table 12.5. Multiple Integrated PC Servers Running Notes					
6x 6506s on CISC (non-EFS)						
	Integrated PC Servers Running Lotus Notes on an AS/400 Model F90					
		V3R2				
	6 IPCS (6506) w/64N	MB RAM, Each on Sep	arate 16Mbps LAN			
				Int PC		
Workload	Number of	AS/400 Model	AS/400 Model	Server (6506)		
Scenario	Users	F90 CPU (%)	F90 IO/Sec	CPU (%)		
Mail	1200 (200 per 6506)	7%	404	59% 61%		
				65% 65%		
				65% 68%		
Note:						
Average utilization of each 16Mbps TRLAN was <1%						
< This environment was not repeated with the IPCS Type 6616>						

IPCS on RISC platforms

The following sections cover results on the RISC platforms. Here are a couple of recommendations before the particulars:

- Use the (RISC) values in this section when projecting performance or capacity of other RISC models. This is warranted due to the differences of the VDASD path and structure on the RISC/SLIC as compared to the same on the CISC/VMC platform.
- Similarly, use the values from the Model 150 tables to project to other Model 150 models. An additional difference is that the Model 150's 2850's Pentium speed is 133mHz compared to the 6616's 166mHz.

Table 12.6. Users on Integrated PC Server Running Notes						
2850 on Model 150 (non-EFS)						
	Lotus Notes	on IPCS Type 2850 w/1	28MB RAM			
	AS/400 Model 150-2	2269 Used, V3R7Enh, 4	DASD Arms; 96MB			
AS/400 AS/400 Int PC						
Workload	Number of	150-2269	150-2269	Server (2850)		
Scenario	Users	CPU (%)	IO/Sec	CPU (%)		
Mail	100	10	38	9		
Mail	250	22	86	23		
Mail	300	29	111	29		
Mail	400	41	161	42		
Mail + Discussion	50	7	25	8		
Mail + Discussion	100	14	51	14		
Note:						
Average utilization of 16Mbps TRLAN was <1% for all scenarios.						
Measurement data on this configuration with EFS is not available.						

IPCS 2850 on RISC(Model 150)

Table 12.7. Users on Integrated PC Server Running Notes 6616 on RISC (non-EFS) Lotus Notes on IPCS Type 6616 w/256MB RAM AS/400 Model 400-2131 Used, V3R7Enh, 12 DASD Arms; 96MB AS/400 AS/400 Int PC Workload Number of 400-2131 400-2131 Server (6616) CPU (%) CPU (%) Scenario Users IO/Sec Mail 100 17 43* 8 200 26 67 15 Mail Mail 300 43 115 24 Mail 400 57 156 36 Mail 490 69 191 54 Mail + Discussion 100 21 55 15 Mail + Discussion 150 31 80 21 Mail + Discussion 200 39 104 27 Mail + Discussion 250 49 131 35 Note: Average utilization of 16Mbps TRLAN was <1% for all scenarios. *I/O rate is higher than expected (38). Measurement data on this configuration with EFS is not available.

Observations

The I/O per second rate on the Model 150/RISC is up to 13% higher than on the CISC model (compared point for point). Part of the difference is the ~5% lower I/O rates (per user) at the 200 and above, user load point on the CISC platform. Except for the RISC's 100 user load point, the other points correlate

fairly well. This consideration contributes to the rationale for using RISC data to project to other RISC and to using CISC data to project to CISC.

Table 12.8. Users on Integrated PC Server Running Notes					
6506 on RISC (non-EFS)					
Lotus Notes on IPCS Type 6506 w/64MB RAM					
	AS/400 Model 400-2	131 Used, V3R7Enh, 12	DASD Arms; 96MB		
AS/400 AS/400 Int PC					
Workload	Number of	400-2131	400-2131	Server (6506)	
Scenario	Users	CPU (%)	IO/Sec	CPU (%)	
Mail + Discussion	60	14	35	37	
Mail + Discussion	100	20	52	57	
Note:					
Average utilization of 16Mbps TRLAN was <1% for all scenarios.					

Observations on the Mail and Discussion Workload

This workload generated a system load that is about 25% higher than the Mail-only load. This is based on the I/O rates listed above in Table 12.7 and Table 12.8.

Client and Server Configurations:

• AS/400 System and IPCS

The AS/400 systems and IPCS used for these measurements were dedicated while executing the workload scenarios described above. Utilization data is provided in the tables for key system resources utilized during the measurements.

Several combinations of IPCS and AS/400 systems were used and are explicitly specified in the headings of the measurement data tables:

- ✤ 6506(64MB RAM) on AS/400 Model E25
- ✤ 6616(256MB RAM) on AS/400 Model E25
- ✤ 6506(64MB RAM) on AS/400 Model 400-2131
- ✤ 6616(256MB RAM) on AS/400 Model 400-2131
- ✤ 6616(256MB RAM) on AS/400 Model 40S-2112
- ✤ 2850(128MB RAM) on AS/400 Model 150-2269
- ♦ 6506(64MB RAM) on AS/400 Model 20S-2010 (w/100mHz clients)
- Multiple 6506(64MB RAM) on AS/400 Model F90 (w/100mHz clients)
- Notes clients Except where noted, all the the PC/Clients used were OS/2 Notes Clients running on Pentium/133mHz hardware.

IPCS with EFS on RISC platforms

The following tables contain results from an IPCS with EFS on a AS/400 RISC model. Note that this is a relatively high speed server system (compared to the other models used in previous measurements). Default EFS cache sizes have been selected to optimize performance in the Notes environment.

For the Mail workload, use of EFS (IPCS with greater than 64MB of RAM):

• reduced DASD I/O request rates

Compared to the previous (non-EFS) measurement data, EFS has lowered I/O rates by 40% to 60%.

• reduced AS/400 CPU consumption

Since the AS/400 CPU is used mainly to service DASD I/O requests, its CPU utilization (or savings) will be proportional to the I/O rates.

• reduced IPCS CPU consumption

IPCS CPU savings at the high end of the workload is about 40%; and no savings at the low end of the workload curve.

• increased capacity/throughput

Capacity of a 256MB IPCS has increased to the 700 level from a previous high of 490 (reached in a non-EFS environment).

Note that the capacity limits reached here are not due to excessive processor(s) and I/O utilizations; but is due to an OS/2 memory management limitation.

For the DB2 Integration environment and workloads, EFS may provide minimal performance improvement. It should however, provide equivalent performance to the non-EFS environment.

Following are data from runs on an IPCS with 256MB, 128MB and 64MB.

Table 12.9. Users on I	Integrated PC Server Runn	ing Notes				
		6616 with EFS on RISC	1			
Lotus Notes on IPCS Type 6616 w/256MB RAM						
	EFS Cache Size Set at 25% (=64MB)					
AS/400 Model 40S-2112 Used, V3R7Enh, 8 DASD Arms; 128MB						
	AS/400 AS/400 Int PC					
Workload	Number of	40S-2112	408-2112	Server (6616)		
Scenario	Users	CPU (%)	IO/Sec	CPU (%)		
Mail	100	1.1	22	9		
Mail	200	1.6	32	13		
Mail	300	2.8	56	20		
Mail	400	4.0	79	25		
Mail	500	5.0	100	34		
Mail	600	6.0	121	40		
Mail	700	8.5	165	49		
Note:						
Average utilization of 16Mbps TRLAN was <1% for all scenarios						

Table 12.10. Users on Integrated PC Server Running Notes										
6616 with EFS on RISC										
Lotus Notes on IPCS Type 6616 w/128MB RAM EFS Cache Size Set at 25% (=32MB) AS/400 Model 40S-2112 Used, V3R7Enh, 8 DASD Arms; 128MB										
								AS/400	AS/400	Int PC
						Workload	Number of	408-2112	40S-2112	Server (6616)
Scenario	Users	CPU (%)	IO/Sec	CPU (%)						
Mail	100	1.1	22	8						
Mail	200	2.1	39	13						
Mail	300	4.0	72	21						
Mail	400	5.3	96	27						
Note:										

Average utilization of 16Mbps TRLAN was <1% for all scenarios

Table 12.11. Users on Integrated PC Server Running Notes					
6616 with EFS on RISC					
Lotus Notes on IPCS Type 6616 w/64MB RAM					
EFS Cache Size Set at 12.5% (=8MB)					
AS/400 Model 40S-2112 Used, V3R7Enh, 8 DASD Arms; 128MB					
		AS/400	AS/400	Int PC	
Workload	Number of	40S-2112	408-2112	Server (6616)	
Scenario	Users	CPU (%)	IO/Sec	CPU (%)	
Mail	100	1.9	33	10	
Mail	200	3.2	57	14	
Note:					
Average utilization of 16Mbps TRLAN was <1% for all scenarios					

12.2 Lotus Notes DB2 Integration Performance

This section contains performance information for three functions provided by Lotus Notes DB2 Integration:(**Note** Currently, there are no measurement or experimental data available for DB2 Integration on IPCS with EFS.) Additional information is in the section "IPCS with EFS on RISC platforms".

• Import

Import refers to the capability to create a Lotus Notes database from existing data stored in a DB2/400 database.

• Shadowing

While performing an import, the shadowing option can be specified which will initiate automatic updates to the imported Lotus Notes database based on on-going changes occurring to the DB2/400 files from which the data was imported.

• Exit Program

Exit program refers to the ability to create and register an exit program that periodically updates DB2/400 data based on changes made to a Lotus Notes database.

Importing DB2/400 data to a Lotus Notes database

In addition to the considerations and recommendations provided in this section, Lotus Notes recommendations for Notes databases should also be reviewed. For example, Lotus Notes recommends that Notes databases be less than 100MB in size. As Notes databases grow to this size and have large numbers of documents, some functions such as initially opening a Notes database can take quite a long time.

The following tables provide measurement data that was collected for imports with varying numbers of rows, sizes of rows, and number of fields per row. In the examples in this section, each imported row from DB2/400 becomes a Notes document. The data was imported from the same AS/400 system as where the IPCS Notes server resides. Very little activity from Notes clients was present during the import measurements described in this section and the AS/400 system was not dedicated but had low resource utilization from other work. The numbers provided in Table 67 below can be used as guidelines to help estimate time to import data into a Notes database from DB2/400.

Table 12.12. Importing DB2/400 Data to a Notes Database					
IPCS 6506 on CISC (non-EFS)					
Importing DB2/400 data to a Notes Database on an IPCS-6506 64MB RAM					
(Background Task AMgr Set to be 0% Active During Imports)					
	AS/400 D60, V3R2, MTU on NWSD = 15400				
Number of					
Imported	Bytes per	Columns per		Resulting	
Rows	Row	Row	Import Time	.NSF File Size	
1000	2150	10	5 min	3 MB	
2000	2150	10	6 min	5 MB	
4000	2150	10	10 min	10 MB	
16000	2150	10	95 min	41 MB	
20000	2150	10	100 min	51 MB	
24000	2150	10	NA	NA	
32000	2150	10	175 min	88 MB	
1000	6233	21	6 min	7 MB	
2000	6233	21	13 min	15 MB	
4000	6233	21	25 min	29 MB	
16000	6233	21	105 min	117 MB	
20000	6233	21	NA	NA	
24000	6233	21	171 min	175 MB	
(background task AMgr set to be 10% active during the following imports)					
16000	6233	21	128 min	117 MB	
20000	6233	21	182 min	146 MB	
24000	6233	21	244 min	175 MB	

Table 12.13. Importing DB2/400 Data to a Notes Database					
IPCS 6616 on RISC (non-EFS)					
Import DB2/400 data to a Notes Database on an IPCS-6616 256MB RAM					
	(Background Task	AMgr Set to be 0% Act	ive During Imports)		
AS	5/400 400-2131, V3R7Ei	nh, 96MB, 12 DASD Arn	ns, MTU on NWSD = 15	400	
Number of					
Imported	Bytes per	Columns per		Resulting	
Rows	Row	Row	Import Time	.NSF File Size	
1000	2150	10	<2 min	2+ MB	
2000	2150	10	2 min	4+ MB	
4000	2150	10	4 min	8+ MB	
16000	2150	10	16 min	33 MB	
20000	2150	10	20 min	41 MB	
24000	2150	10	23 min	49 MB	
32000	2150	10	31 min	65 MB	
1000	6233	21	3 min	7 MB	
2000	6233	21	5 min	13 MB	
4000	6233	21	9 min	27 MB	
16000	6233	21	37 min	106 MB	
20000	6233	21	46 min	132 MB	
24000	6233	21	57 min	160 MB	
Note:					

The size than in Table 67 Reasons unknown

Observations

Great performance improvements have been achieved in the DB2/400 Import function with the 6616 IPCS!

- with 2150 bytes/row database, we get up to 6x response time improvement (16,000 imported rows) ٠
- ٠ Import rate is about .4 MB/minute vs 2 MB/minute
- with 6233 bytes/row database, we getup to 3x response time improvement (24,000 imported rows) ٠
- Import rate is about 1 MB/minute vs 2.8 MB/minute •

Caution The Import function may consume substantial resources and could impact other foreground jobs! Note the utilizations and rates as follows during the Import of 16,000 rows (6233 bytes/row) to a Notes database.

- RISC CPU% at 46%
- I/O per second at 130 ٠
- **IPCS CPU%** at 25%

Recommendations

Maximum Transfer Unit

It is recommended to set the Maximum Transfer Unit (MTU) size for the internal LAN to 15400 in order to provide maximum data throughput when using the DB2 Integration functions. The MTU parameter is found on the Network Server Description (NWSD).

AMgr settings

The AMgr background task affects the performance of import function when it is active. It was
observed that with default setting of 50% active daytime and 70% nighttime, that the time to import can vary greatly. These settings can be adjusted and it is recommended for optimum import performance that the setting for percent active time be set near 0 for the time of day when the imports will be occurring. It is also possible to not have this task start up at all by editing the NOTES.INI file and taking it out of the list of tasks that get started on the Notes server. Please refer to the Lotus Notes 4.0 Administrator's Guide for detailed information. The AMgr (Agent Manager) task allows control of who can run agents and when they can run on each server. From the data in Table 67 on page 165 it can be seen that with the AMgr task 10% active that Import times were greater than when AMgr was set to 0% active.

• Starting an import

Import requests are queued and processed one at a time in the order they are requested. If a combination of large and small imports are to be requested, it may be appropriate to submit requests for the smaller imports first (for example if Notes clients are waiting to use them), rather than delaying them the length of time of the larger imports. An agent program wakes up every 2 minutes on the Integrated PC Server to check for import requests, so import requests may take 2 minutes before they begin processing.

• Opening a database

After importing a database the default view is to show all of the columns. Choosing fewer columns through a view can significantly improve the time to do the first open. Notes databases can takes minutes to open the first time if they contain tens of thousands of documents. Refer to Notes on-line help documentation for suggestions on improving view display times.

• Import impacts to Notes clients

While an import is active, the performance of Notes clients can be impacted. When possible it is recommended to perform imports when the level of Notes client activity is low.

• Import versus Import with Shadowing consideration

Importing a given number of rows will take somewhat longer if the import with shadow option is specified. Please refer to data in the next section on shadowing for additional information.

Shadowing DB2/400 data to a Lotus Notes database

When initiating an import, the user is given an option to start shadowing function for the DB2/400 database file(s) indicated in the import request. To shadow DB2/400 files to a Notes database, Data Propagator Relational/400 is used and the files to be imported from and shadowed must have journaling active. For further information on shadowing, please refer to the resources indicated at the beginning of this section as well as **DataPropagator Relational Capture and Apply/400** (SC41-3346) and. Other parameters that can be specified when initiating an import with shadowing include the time of day and/or frequency at which the user desires the shadowing activity to occur. Considerations for these setting are discussed later in this section.

The data in the following tables provide examples of the time required to shadow various types of changes (inserts, deletes, updates) from a DB2/400 file to a Notes database.

Table 12.14. Shadowing DB2/400 Data to Notes Database						
IPCS 6506 on CISC (non	-EFS)					
Shadowing DB2/400 Data to a Notes Database on IPCS-6506 64MB RAM						
(Background Task Set to be 0% Active D	uring Measurements)					
Documents in Notes Database We	ere 2150 bytes					
AS/400 D60 V3R2, MTU on NV	VSD = 15400					
Description of Shadowing Changes Made to 10,000	Time to Shadow Changes					
Document Notes Database						
200 Inserts	20 minutes					
200 Deletes (spread throughout Notes DB)	12 minutes					
200 Updates (spread throughout Notes DB)	11 minutes					
100 Inserts, 100 Deletes, 200 Updates	27 minutes					
Description of Shadowing Changes Made to 10,000	Time to Shadow Changes					
Document Notes Database						
200 Inserts	49 minutes					
100 Deletes (spread throughout Notes DB)	14 minutes					
200 Updates (spread throughout Notes DB) 24 minutes						
100 Inserts, 100 Deletes, 200 Updates	54 minutes					
Note:						

The "time to shadow changes" only includes the time to shadow the changes to the Notes database, the DB2/400 file changes had already occurred.

Table 12.15. Shadowing DB2/400 Data to Notes Database									
IPCS 6616 on RISC (non-EFS)									
Shadowing DB2/400 Data to a Notes Database on IPCS-6616 256MB RAM									
(Background Task Set to be 0% Active During Measurements)									
Documents in Notes Database Were 2150 bytes									
AS/400 400-2131, V3R7Enh, 12 DASD Arms, 96MB;	MTU on NWSD = 15400								
Description of Shadowing Changes Made to 10,000	Time to Shadow Changes								
Document Notes Database									
200 Inserts	3 minutes								
200 Deletes (spread throughout Notes DB) 2 minutes									
200 Updates (spread throughout Notes DB) 2 minutes									
100 Inserts, 100 Deletes, 200 Updates	4 minutes								
Description of Shadowing Changes Made to 10,000	Time to Shadow Changes								
Document Notes Database									
200 Inserts	6 minutes								
100 Deletes (spread throughout Notes DB)	2 minutes								
200 Updates (spread throughout Notes DB)	3 minutes								
100 Inserts, 100 Deletes, 200 Updates 7 minutes									
Note:									
The "time to shadow changes" only includes the time to shadow the changes to the Notes database, the DB2/400 file									
changes had already occurred.									

Observations

This set of Inserts, Deletes, Updates test cases also exhibited substantial improvements in response times averaging

- 6x improvement when operating on the 10,000 document Notes database.
- 8x improvement when operating on the 20,000 document Notes database.

Conclusions/Recommendations:

• Inserts

Shadowing inserts to Notes databases occurs faster for smaller databases and proportionately longer for larger Notes databases. From the data in Table 69 on page 168

- with 6506/CISC, inserts into the 10,000 document Notes database occurred at a rate of approximately 600 per hour, while inserts into the 20,000 document database occurred at a rate of approximately 250 per hour.
- ♦ with 6616/RISC, it'd be 4000/hour and 2000 per hour respectively.
- Deletes and Updates

From Table 12.04.

✤ On the 6506/CISC platform:

when performing Shadowing, Deletes and Updates on the 10,000 document (Notes database); the deletes and updates occurred at a rate of approximately 1000 per hour, and at a rate of approximately 400 to 500 per hour for the 20,000 document Notes database.

Similar to insert, the size of the Notes database will typically impact the rate at which shadowed deletes and updates can occur. For these examples the documents to be deleted or updated were spread throughout the Notes database. If the documents had been found at the top of the Notes databases in both cases, the rates for shadowing delete and update would have been similar for the 10,000 and 20,000 document databases.

- on the 6616/RISC platform, the rates are 6000/hour and 3600/hour respectively.
- Shadowing vs Importing

From the data in Table 12.04168 above, it can be seen that the various types of changes can be shadowed at the rate of hundreds of changes per hour. The entire 10,000 document database that was being updated in the first part of Table 168 was imported in 32 minutes. Many issues need to be considered regarding the use of the Notes database, but if the changed data is not required by applications or users in a real time manner, it may be a consideration to repeatedly import the entire database if thousands of changes are anticipated for the DB2/400 files that would be shadowed.

• Shadowing impact on Notes clients

Like import, shadowing activity can significantly impact the performance of active Notes clients. When initiating the import with shadowing request, consider the settings for frequency of shadowing interval. This interval determines how often the Notes server checks to see whether shadowing changes are queued up and then begins making the changes to the Notes database. To avoid impact to Notes clients, attempt to schedule shadowing activity during times when client activity is low.

• Shadowing impact on Import

Import with Shadowing option ON will take longer than Import without Shadowing. From data in Table 12.02165, a 20,000 row (2150 bytes per row) import took 100 minutes. Importing the same 20,000 rows with the shadowing option specified took 132 minutes (data for time to import with shadowing is not shown).

• Estimating shadowing rates

Data from Table 168 can be used to estimate time to complete various combinations of shadowing activity. Using the individual rates for inserts, updates, and deletes of a 10,000 document database to estimate the time to perform 100 inserts, 100 deletes, and 200 updates would yield an estimate of 27 minutes. This is in fact how long it actually took when that specific combination of changes was measured as indicated in the table. Data in Table 70 on page 168 can be used in a similar manner.

Exit Program: data from a Notes database to DB2/400

It is up to the user to create and register the Exit program to process the data once it is stored in the AS/400. The data in Table 12.16 provides examples for times it took to send the changes from the Notes database to DB2/400 for update. This environment used a 10,000 document Notes database and the rates of changes for inserts, updates and deletes are very similar to those shown in Table 12.14 for Shadowing changes.

Table 12.16. Exit Program Function						
IPCS 6506 on CISC (non-EFS)						
Exit Program: Changed Data From Notes Database to DB2/400, Does Not						
Include Processing to Apply Changes to DE	32/400					
(Background Task Set to be 0% Active During Me	easurements)					
Documents in Notes Database Were 2150	bytes					
AS/400 D60 V3R2, IPCS-6506 64MB RAM, MTU on	NWSD = 15400					
<this 2850="" 6616="" environment="" ipcs="" not="" or="" repeated="" the="" type="" was="" with=""></this>						
Description of Changes Sent to DB2/400 Based on	Time to Shadow Changes					
Changes to Notes Database						
200 Inserts 22 minutes						
200 Deletes (spread throughout Notes DB) 12 minutes						
200 Updates (spread throughout Notes DB) 13 minutes						
100 Inserts, 100 Deletes, 200 Updates	29 minutes					

Chapter 13. Language Performance

This chapter focuses on Language Performance for languages other than Java. For Java-related information, refer to the chapter on Java Performance found in this document.

From its inception, the AS/400 was designed to easily incorporate new technology while protecting customers' substantial investments in software. AS/400 Advanced Application Architecture effectively isolates applications from low level changes in the system. Changes to accommodate new technology are made below the Technology Independent Machine Interface, which is preserved across releases. As a result, applications running above the Technology Independent Machine Interface do not have to change.

The program creation process on AS/400 is one of the key elements in isolating applications from low level changes in the system. With the introduction of AS/400 RISC models, there have been changes to this support that directly affect compile times. To discuss these changes and the significant improvements in ILE compile times, this section will cover the following topics:

- How programs are created, and what changed for AS/400 RISC models
- Compile time performance
- Runtime performance
- Program object size comparisons
- Working memory guidelines for compiles

Note: The performance comparisons mentioned in this chapter apply only to software release-to-release comparisons. Each hardware release also provides an additional level of performance improvement.

CISC Program Model

When you compile an AS/400 program, the system goes through a two step process (see Figure 13.1 below). The system first creates a program creation template (PCT). The PCT is independent of the instruction set that is used by the processor. The system then translates the PCT into a set of machine instructions. The end result is an executable program (*PGM) or an ILE module (*MODULE). Under OS/400 V2R3 and later CISC-model releases, two system translators are used, one for OPM languages and one for ILE languages.



Figure 13.1

The PCT is stored with the program unless you remove the program's observability. The AS/400 Advanced Application Architecture allows you to translate the same PCT into RISC instructions or CISC instructions. With this capability, you can move programs between RISC and CISC platforms without recompiling. As discussed earlier, this is a major advantage of the AS/400 as it allows you to incorporate the latest technology without having to rewrite or recompile your application.

RISC Program Model

To get the full performance benefit of the AS/400 PowerPC RISC processor architecture, an advanced code generation technology is required. Code optimization and instruction scheduling increases your application's performance by eliminating redundant instructions and reducing unused cycles. This advanced technology is part of the Optimizing Translator.

To make this technology available to the existing OPM compilers, the OPM PCTs are automatically transformed into ILE PCTs. These ILE PCTs are then used by the Optimizing Translator to generate optimized RISC instructions (Figure 13.2 below). With this you get the benefits of advanced code generation and application investment protection automatically. The additional conversion step does affect the compile time performance of OPM programs which is discussed in the next section.





13.1 Compile Time Performance

The purpose of this section is to provide you with general information on compile time performance when moving to an AS/400 RISC model. First, here's a list of system changes that significantly impact compile time. The list begins with the two factors mentioned in the preceding section:

- RISC-model advanced code generation
- OPM program creation template conversion
- Memory requirements

Note: These conclusions are drawn from a set of measurements that were made during the early releases of RISC (V3R6 and V3R7). Although they haven't been remeasured on the current release, it is expected that compile-time performance has slightly improved or remained unchanged with each software release. It

is nearly impossible to make a general performance statement that applies to all applications and systems, due to the variety of factors that affect compile-time performance. Most environments will see better performance; however, some might see degradations. These conclusions are meant to be a general guide and are not meant to suggest minimum performance expectations, or to guarantee performance for any particular application.

Compile Time Conclusions when comparing CISC and RISC

The compile process requires more memory on RISC models than CISC models. The following conclusions are based on using sufficient memory to keep paging to a minimum. Compile times are also sensitive to the optimization level.

- The compile time of ILE languages has improved significantly on RISC models. The RISC-model ILE compilers are typically 2 times faster than their CISC-model counterparts when compiling with optimization level (*NONE). When compiling with optimization level *FULL, the RISC-model compile times are generally 30-40% faster.
- The most improvement in compile time will be compiling large programs at optimization level (*NONE). The least improvement will be compiling small programs at higher optimization levels.
- OPM compiles on RISC releases are approximately 20-70% slower than OPM compilers on CISC releases for non-optimized programs. The additional time is the result of the automatic program creation template conversion to allow OPM languages to utilize the Optimizing Translator.
- OPM optimized compiles on RISC releases are much longer than on CISC releases. This is due to the more advanced optimizations performed by the Optimizing Translator when compared with the CISC-model OPM translator.

Compile Time Conclusions for later RISC releases

In follow-on RISC releases the following changes were made that could affect compile time:

- In the V3R7 release, there was a significant reduction in compiler working set size, resulting in a considerable improvement for compiles in a memory-constrained environment.
- In V3R7, a change was made in the register allocation strategy, resulting in a significant improvement to runtime code. In some cases, this change caused compile-time degradations for large compiles (modules with procedures larger than 500k bytes) at OPT(*FULL) or above. These compiles may take twice as long as they did in V3R6. If the compile time is too long, the module should be restructured with smaller procedures or the compiler optimization level reduced to OPT(*BASIC).

<u>Compile Time Recommendations</u>

• When possible, application developers should move to ILE. In addition to improved compile times, ILE offers many advantages over OPM, such as modularity, static binding, common run-time services, and improved code optimization.

ILE RPG is shipped with a command, CVTRPGSRC, which can be used to migrate your RPG III source code to RPG IV. The *ILE RPG/400 Programmers Guide*, *Appendix B*, contains a detailed

description of the conversion process, with examples to help you identify and quickly resolve potential conversion problems. Another source for information on the conversion process is the Redbook entitled *Moving to ILE RPG (GG24-4358)*.

For conversion and compatibility considerations between OPM COBOL and ILE COBOL for OS/400, please refer to the *ILE COBOL/400 Programmers Guide, Appendix G*.

- The following suggestions help in managing and improving compile times:
 - ♦ For initial compiles use OPTION(*NOGEN), and optimization *NONE or *NOOPTIMIZE.

*NOGEN compiles the module or program but does not create a program object. It can be used to fix and edit compile errors.

Using optimization *NONE or *NOOPTIMIZE can dramatically reduce compile times. Optimized compiles can be expected to take at least 3-5 times longer than compiles at optimization *NONE or *NOOPTIMIZE. Once the application is debugged and ready for production use, compile it at the appropriate optimization level, and conduct a final test. Typically RPG and COBOL programs should be compiled at optimization level (*NONE) or (*BASIC), and C and C++ programs at optimization level (*FULL) or level 40.

✤ Use the appropriate working memory size

See Working Memory Size Guidelines section below.

- ✤ Compile in batch rather than interactively.
- The following recommendations hold for ILE applications:
 - Design modular applications

Modular programming offers faster application development and a better ability to reuse code. Programs are developed in smaller, more self-contained procedures. These procedures are coded as separate functions, and then bound together to build an application. By building applications that combine smaller and less complex components, you can compile and maintain your programs faster and easier.

Use the value of DBGVIEW adequate for your purpose.

Requesting debug information requires more compile time and creates larger objects. For example, DBGVIEW(*LIST) results in a slower compilation time than DBGVIEW(*STMT). If the level of debug information you need is that provided by DBGVIEW(*STMT), selecting *LIST would unnecessarily slow down compilation time and inflate object size.

13.2 Runtime Performance

Note: These conclusions are drawn from a set of measurements that were made during the early releases of RISC (V3R6 and V3R7). Although they haven't been remeasured on the current release, it is expected that runtime performance has improved with each software release. The level of performance improvement is highly dependent on the characteristics of the application being measured and the type of system services used by that application. It is nearly impossible to make a general performance of an application, as well as the numerous system configurations that are available. Most environments will see better performance. These conclusions are meant to be a general guide and are not meant to suggest minimum performance expectations, or to guarantee performance for any particular application.

Runtime Conclusions

- Record I/O functions, which are faster than stream I/O, don't show as much relative gain over CISC releases as do other C functions. This is because more of the CPU time is spent in OS/400 and database than for the other benchmarks.
- General logic and integer computation for RISC releases are much improved over that of CISC releases; generally between 2X and 3X faster. Floating point applications gain less than integer based applications.
- Call intensive functions are much faster on RISC releases than on CISC releases. At OPT(30) gains are about 3X. At OPT(40), performance of leaf functions (those functions not calling other functions), is particularly better, averaging about 5X faster.
- In the first RISC release, V3R6, a change was made to do a direct mapping of stream I/O interface to the Integrated File System (POSIX) APIs. Optimum stream I/O performance is achieved when data is stored in the QOpenSys file system. Early measurements have indicated that stream I/O performance can improve significantly when using IFS.

Tradeoffs

- At OPT(10), average C applications on RISC systems gain about 2X relative to CISC systems.
- OPT(20) yields about 10% performance improvement over OPT(10), at a cost of increased compile time (1.7X slower than OPT(10)).
- OPT(30) yields about 45% performance improvement over OPT(10), at a cost of increased compile time (4X slower than OPT(10)).
- OPT(40) yields about 5% performance improvement when compared to OPT(30). The compile time is about the same as OPT(30). This improvement depends on how many procedure calls are to leaf routines. Some programs may see a larger benefit from OPT(40).

Run Time Recommendations

To take full advantage of the latest compiler optimization technology that is available on AS/400 RISC models, it is recommended that application developers move to ILE. In addition to improved compile times, ILE offers many advantages over OPM, such as modularity, static binding, common run-time services, and improved code optimization. Refer to the section entitled "ILE Compiler Optimization" for more details and reasons why ILE is better for RISC.

13.3 Program Object Size Comparisons

Program Object Size Growth

On RISC releases, there are several architectural factors which influence the size of program objects. First, the page size has increased from 512 bytes to 4K bytes (4KB). The larger page size is important in making storage management algorithms more efficient as the size of main storage continues to increase. However, the 4KB page does impact the size of objects, particularly smaller objects, since objects must be aligned on 4KB boundaries rather than 512 byte boundaries.

Second, in general, the number of instructions for a comparable program is going to be larger on RISC systems than on CISC systems. This is referred to as code expansion. By the very nature of RISC design (efficient execution of simple instructions), it takes more instructions to do the same function as on CISC. For example, on RISC there are no storage-to-storage instructions; all data must be processed through registers. On CISC, moving data between two storage locations can be done with a single MVC (move character instruction). On RISC this requires a Load and then a Store instruction. In most cases, code expansion will have more of an impact on program object growth than the 4KB page size.

Conclusions and Recommendations

Due to the above factors, program object growth when moving to a RISC system from a CISC system is as follows:

- Observable ILE programs will grow on the average by 1-2 times.
- Observable OPM programs will grow on the average by 2-3 times.
- Non-observable OPM programs will grow on the average by 4 to 5 times. This range may be as high as 6 times for very small programs and as low as 3 times for very large programs.

The relative growth of non-observable OPM programs on a RISC system when compared to a CISC system will be greater than the relative growth for observable OPM programs because the majority of program growth is due to the code expansion of the executable part of a program object. The size of the program creation template (PCT) does not increase significantly, and for non-observable programs, the PCT has been removed. With the exception of observable programs with no compression, the ILE versions of programs are slightly smaller than the OPM versions. Also, the effectiveness of optimization as a size reduction tool is not large. A reasonable expectation would be that individual programs are reduced by one 4KB page per level of optimization.

To determine overall system DASD growth when moving from CISC to RISC, you should use the Upgrade Assistant. For more information on the Upgrade Assistant refer to the manual *AS/400 Planning for PowerPC Technology, SA41-4154*.

Here are several options to consider to reduce storage requirements for your program objects:

• Compress the observable part of the program

After you build your application, before it goes into production, you can reduce all unused parts of the object until they are actually needed by using the CPROBJ (Compress Object) command. Specify *OBS on the PGMOPT parameter. You should NOT remove program observability to reduce program size.

- For ILE programs, use modular design techniques and select the correct level of debug options when compiling your program.
 - Using service programs as a means of reusing code reduces the overall storage requirements for your application.
 - Generating DBGVIEW data may increase program object size significantly depending on the DBGVIEW options used.
 - If DBGVIEW(*LIST) compile option is chosen, the compile listing used for debugging is stored with the object, thus greatly increasing the program object. Carefully weigh the advantage of having a compiler listing stored with your object against the additional storage requirements.
 - Consider using DGBVIEW(*SOURCE). It may give you similar capabilities in debug, but results in a smaller program object size.

13.4 Working Memory Guidelines for Compiles

Working memory size is the amount of memory required to do a task satisfactorily. Think of working memory size this way: Given infinite memory, the compiler will run at its optimal speed. If you restrict memory, the compiler will have to swap pages to DASD making it run slower. The more memory is restricted, the more time the compiler spends swapping memory pages. As an example, look at a program with 1500 C specifications and 5300 MI instructions. This program would compile reasonably quickly using an 8 MB pool if the program was not optimized. The optimized program's compilation will benefit from as much memory as you can give it, although there is not much benefit beyond 64 MB.

Conclusions/Recommendations

• As a general rule, to achieve minimum OPM and ILE compile times, use 16-20 MB to compile a medium-size program and 32MB to compile a large program. Smaller pool sizes will result in longer compile times, and less than 8MB is not recommended.

- Regarding minimum system configurations, a system with 32 MB of memory may be sufficient for casual application development work (for example infrequent compilations). This of course also depends on what other workload is running on the system. For systems that are used primarily for application development work, a minimum of 64MB of memory should be considered.
- As a general guideline for the memory size of systems used primarily for application development work, you should assume each concurrent compile requires about 25MB of main storage. For example, if the system needs to support 10 concurrent compiles, then as an initial estimate, the memory size of the system should be 256MB. If there is other work in addition to the application development work, the main storage requirements for that work needs to be taken into account also. For detail system capacity sizing, you should use BEST/1 for OS/400. BEST/1 takes into the account the additional main storage required for application development workloads, and should be used to accurately size main storage needs.

13.5 Application Compute Intensive Performance

In general, the performance improvement of applications referred to as 'application compute intensive' is significantly more than the improvement of traditional commercial applications when moving from CISC to RISC technology.

This section will cover:

- Performance of traditional commercial applications
- What is meant by 'application compute intensive'
- Why the performance of these application types improve significantly on RISC technology

Traditional Commercial Applications

If you look at the CPU time profile for a traditional commercial application, typically up to 10-20% of the CPU time is spent in application programs, while the remaining 80+% is spent in operating system programs (Figure 37 below). This is because traditional commercial applications spend much of their CPU time in system services such as database I/O, query processing, workstation/printer processing, and communication I/O.

Since a large part of the CPU time for commercial applications is spent in system services, performance improvements in these applications would result from either changing the application to utilize system services more efficiently, or from performance improvements to the system services. The performance of AS/400 system services has been optimized over several releases of OS/400 and perform equivalently on RISC as compared to CISC. As a result, the performance of traditional commercial applications on RISC systems will be equivalent to CISC systems with the same relative performance rating. For these types of applications, RISC offers improved price/performance and significant potential for performance growth over CISC with the introduction of PowerPC technology.

Application Compute Intensive

As compared to the traditional commercial applications, there are AS/400 applications where much more of the CPU time is spent in application programs (Figure 13.3 below). These types of programs are referred to as *Application Compute Intensive*. For example, applications that implement complex business rules for decision making are typically application compute intensive, as are financial modelling applications that do a significant number of numeric calculations.

Another example of applications that may be application compute intensive are the growing number of portable applications available on the AS/400. To achieve a high level of portability, these applications typically use only functions widely available on a number of systems. They implement their own functions which are provided more efficiently by unique OS/400 system services. As a result more CPU time is spent in application code.

The amount of performance improvement of applications that are application compute intensive depends on the actual workload, but you can expect to see improvements ranging up to 3 times when moving to RISC.





ILE Compiler Optimization

Although applications that are application compute intensive can be written in other languages, they are typically written in ILE C and C++. The main reasons for the significant performance improvements of these kinds of applications on RISC are:

- The nature of compute intensive applications in C and C++ are fundamentally different than that of traditional commercial applications. For example, they typically have more loop iterations, use pointers more intensively, and use integer as compared to decimal data types. These types of operations fit the RISC computation model more closely than code generated for commercial applications.
- C and C++ coding paradigms give ILE more opportunities for optimization, such as strength reduction of loops, and common subexpression elimination.
- ILE can better exploit the power of RISC hardware. For example, the superscalar design of RISC provides multiple instruction pipelines which permit multiple instructions to be executed at the same time. The Optimizing Translator takes advantage of the superscalar design using instruction scheduling to resequences instructions to maximize instruction overlap. This was not (generally) available on CISC.

13.6 Conclusions and Recommendations

There are many functional and performance advantages in moving to ILE. For more details refer to *ILE Concepts SC41-5606*. This reference describes binding considerations, activation groups, as well as the advanced optimizations that are available.

For language specific performance tips and techniques, refer to the appropriate language's ILE Programmer's Guide. For example, the *ILE C Programmer's Guide SC09-2712* gives many coding tips for improving performance for applications written in ILE C. This is especially important for Application Compute Intensive applications.

For general AS/400 application performance tips and techniques refer to the AS/400 Performance Management V3R6/V3R7 Redbook SG24-4735-00.

For information on Java Performance, refer to the chapter on Java Performance found in this document.

Chapter 14. DASD Performance

The change from a 512 byte to a 4KB page size will in general not noticeably change the DASD response time characteristics of an application as long as sufficient memory is added when upgrading to RISC. This is particularly true for batch applications that typically have I/O sizes that exceed 4KB. For applications such as interactive applications doing mostly random I/O, the 4 KB page size may decrease the number of I/O operations, depending on if additional data is accessed in the same 4KB page.

**NOTE*: The 4KB page size plus code expansion due to the RISC architecture will result in increased main storage requirements over IMPI. Refer to 19.3, "Main Storage Sizing Guidelines" for a discussion on how much main storage is required on RISC as compared to IMPI.

14.1 Device Performance Characteristics

This section compares the performance of the Internal DASD Subsystems based on the 65x2 RAID Controllers or 6530 Storage Controller with the external 9337 Disk Array Subsystem using a system configured with an equivalent amount of DASD capacity. This section also contains performance characteristics for the 6532, 6533, 2726, 2740, 2741 and new 2748 RAID Controllers and 6751, 6754 and 9728 MFIOPs. The performance is based on measurements and modeling done in the development laboratory. Because the performance of the AS/400 system is dependent on many factors, these characteristics are very general in nature. To assess the various configuration options, one of the capacity planning tools should be used.

The performance characteristics of Internal DASD is listed in Table 14.1 and Table 14.2 below and the performance characteristics of External DASD is listed in Table 14.3. One new DASD model (10K RPM 6717 - 9GB capacity) has been announced for V4R4. The tables do not list all of the feature codes, but it does provide performance information for most of the disk configurations. For example, the 6522 IOP has the same performance characteristics as the 6502 IOP. For a description of the DASD models supported by the 6502,6512, 6530, 6532, 6533, 6751, 6754, 2726, 2740, 2741, 2748 and 9728 IOP/IOAs refer to Appendix C, "DASD IOP Device Characteristics".

In the tables, the following measures of performance are listed.

Service Time is the time required to perform the "Interactive op" described in the next paragraph. The time starts with the request from the CPU to the Disk IOP and the time stops when the data is in main storage (read) or when the data is on the disk or in the write cache (write). Queueing time is not included.

Interactive Ops/Sec is an estimate of the number of IOs that can be done at 40% utilization using the service time calculated for the previous column. If the disk model contains 2 arms, this number only reflects the capacity of one arm. At 40%-50% utilization, the disk arms are at the "knee of the curve". As utilization exceeds the "knee of the curve", response time increases significantly and becomes erratic. We assume the following:

- 40% arm utilization
- 7KB transfer size
- 70% read and 30% write
- 80% 1/3 seek and 20% 0 seek

Interactive Rel is the Relative Interactive performance of the disk drives. This column is the same as the INTERACTIVE Ops/Sec column except that the numbers are normalized to 1.0.

Batch Hours is an estimate of how long batch type applications would execute. The duration of many batch type jobs depends on the performance of the disk. For ease of understanding, the numbers are normalized to 8 hours assuming the slowest disk drive is used. We assume the following:

- 75% of the batch job time is disk IO
- Average of 4KB, 8KB and 16KB transfer sizes
- 70% read and 30% write
- 20% 1/3 seek and 80% 0 seek

Ops/Sec/GB is an estimate of how many system physical disk IOs per second per usable GB of space that the specific model of DASD can perform when the arm is 40% utilized. The write cache effectiveness reduces the volume of writes that the physical disk drive must support. For the 9337-2xx models, the write cache effectiveness is assumed to be 45% and for the 9337-4xx and 9337-5xx models it is assumed to be 65%. For the 6502 IOP, the write cache effectiveness is assumed to be 55%.For the 6512, 6532, 6533, 6751, 6754, 2726, 2740 and 2741 IOP/IOAs, the write cache effectiveness is assumed to be 65%. We use the service time required to physically write the record to DASD. The service time contained in column four included the faster write completions that resulted when the write was safely in the write cache.

Table 14.1. DASD Per	formance - In	nternal DASL)					
Disk		Number	Service	Interactive		Batch Ops/Sec/		ec/GB
Model	MB	of Arms	Time	Ops/Sec	Rel	Hours	Base	RAID
6502-6605	1031	1	8.7	46.0	2.3	3.6*	38	29
6502-6606	1967	1	8.8	45.5	2.3	3.6*	20	15
6502-6607	4194	1	8.8	45.5	2.3	3.6*	9	7
6502-6713	8589	1	9.1	44.0	2.2	3.6*	4	3
6502-6714	17548	1	9.1	44.0	2.2	3.6*	2	2
6512-6605	1031	1	8.2	48.8	2.4	3.4*	43	36
6512-6606	1967	1	8.3	48.2	2.4	3.4*	22	18
6512-6607	4194	1	8.3	48.2	2.4	3.4*	10	8
6512-6713	8589	1	8.6	46.5	2.3	3.4*	5	4
6512-6714	17548	1	8.6	46.5	2.3	3.4*	2	2
6530-6605	1031	1	11.4	35.1	1.7	3.9	34	
6530-6606	1967	1	11.6	34.5	1.7	3.9	18	
6530-6607	4194	1	11.6	34.5	1.7	3.9	8	
6530-6713	8589	1	12.0	33.3	1.7	4.0	4	
6530-6714	17548	1	12.0	33.3	1.7	4.0	2	

Note:

The 6502 and 6512 IOP write cache is only used for 1 GB and larger DASD. The write cache is NOT used for any 400 MB or 988 MB DASD that are attached.

* For the 6502 and 6512 IOPs in RAID mode, most batch jobs will run nearly as fast as if they were run in 'base' or mirrored mode. Only in extreme cases will the RAID mode cause degradation. An example is when there are sequences of hundreds of writes to a single IOP in a short period of time.

Table 14.2. DASD Performance - Internal DASD (continued)								
Disk		Number	Service	Interactive Batch		Batch	Ops/Sec/GB	
Model	MB	of Arms	Time	Ops/Sec	Rel	Hours	Base	RAID
6533-6605	1031	1	8.0	50.0	2.5	3.4*	43	36
6533-6606	1967	1	8.1	49.4	2.5	3.4*	23	19
6533-6607	4194	1	8.1	49.4	2.5	3.4*	11	9
6533-6713	8589	1	8.4	47.6	2.4	3.4*	5	4
6533-6717	8589	1	7.2	55.6	2.8	3.1*	6	5
6533-6714	17548	1	8.4	47.6	2.4	3.4*	2	2
2741-6605	1031	1	8.0	50.0	2.5	3.4*	43	36
2741-6606	1967	1	8.1	49.4	2.5	3.4*	23	19
2741-6607	4194	1	8.1	49.4	2.5	3.4*	11	9
2741-6713	8589	1	8.4	47.6	2.4	3.4*	5	4
2741-6717	8589	1	7.2	55.6	2.8	3.1*	6	5
2741-6714	17548	1	8.4	47.6	2.4	3.4*	2	2
2748-6607	4194	1	7.2	55.6	2.8	3.1*	12	10
2748-6713	8589	1	7.5	53.3	2.7	3.2*	6	5
2748-6717	8589	1	6.5	61.5	3.1	3.0*	7	5
2748-6714	17548	1	7.2	55.6	2.8	3.1*	3	2
6754-6605	1031	1	8.0	50.0	2.5	3.4*	43	36
6754-6606	1967	1	8.1	49.4	2.5	3.4*	23	19
6754-6607	4194	1	8.1	49.4	2.5	3.4*	11	9
6754-6713	8589	1	8.4	47.6	2.4	3.4*	5	4
6754-6717	8589	1	7.2	55.6	2.8	3.1*	6	5
6754-6714	17548	1	8.4	47.6	2.4	3.4*	2	2
9728-6605	1031	1	10.9	36.7	1.8	3.8	38	
9728-6606	1967	1	11.0	36.4	1.8	3.8	18	
9728-6607	4194	1	11.0	36.4	1.8	3.8	9	
9728-6713	8589	1	11.3	35.4	1.8	3.9	4	
9728-6717	8589	1	9.7	41.2	2.1	3.6	5	
9728-6714	17548	1	11.3	35.4	1.8	3.9	2	

Note:

* The 6533 IOP has slightly better performance than the 6532 IOP but will usually be noticeable only at higher throughput ranges. The 2741 IOA has slightly better performance than the 2726 IOA but will usually be noticeable only at higher throughput ranges. The 6754 MFIOP has the same performance relationship with the 6751 MFIOP. The *2740 IOA* (which is targeted for smaller systems) has similar performance to the 2726 IOA over typical operating ranges, but has slightly slower performance at higher throughput ranges.

* For the 6532, 6533, 6754, 2726, 2740, 2741 and new 2748 IOP/IOAs in RAID mode, most batch jobs will run nearly as fast as if they were run in 'base' or mirrored mode. Only in extreme cases will the RAID mode cause degradation. An example is when there are sequences of hundreds of writes to a single IOP in a short period of time.

These IOP/IOAs are also capable of attaching Ultra-SCSI (40 MB/sec bus) versions of the 6606, 6607, 6713, and 6714 DASDs. These devices can improve performance for workloads characterized by large disk I/O operations. The new 2748 IOA is capable of supporting the SCSI Wide-Ultra2 (80 MB/sec) bus.

Table 14.3. DASD Pe	rformance - E	External DAS	D					
Disk		Number	Number Service Interactive		Batch	Ops/S	ec/GB	
Model	MB	of Arms	Time	Ops/Sec	Rel	Hours	Base	HA
9337-210	1084	2	12.4	32.3	1.6	4.3*	50	36
9337-215	1084	2	9.8	40.8	2.0	3.9*	63	46
9337-220	1940	2	12.5	32.0	1.6	4.3*	28	20
9337-225	1940	2	11.0	36.4	1.8	4.0*	32	23
9337-240	7868	4	11.3	35.4	1.8	4.0*	15	11
9337-420	3880	4	8.6	46.5	2.3	3.5*	43	36

9337-440	7868	4	8.8	45.5	2.3	3.5*	21	18
9337-480	16776	4	9.1	44.0	2.2	3.6*	10	8
9337-540	7868	4	8.6	46.5	2.3	3.5*	21	18
9337-580	16776	4	8.6	46.5	2.3	3.5*	10	8
9337-590	34356	4	8.9	44.9	2.2	3.6*	5	4
Note:								

* For the 9337-2xx, 9337-4xx, and 9337-5xx models in HA mode, most batch jobs will run nearly as fast as if they were run in 'base' or mirrored mode. Only in extreme cases will the RAID mode cause degradation. An example is when there are sequences of hundreds of writes to a single IOP in a short period of time.

Conclusions / Recommendations

The 6532, 6751 and 2726 DASD IOP/IOAs have similar performance characteristics. The 6533 IOP, 2741 IOA and 6754 MFIOP have slightly better performance characteristics, which are more beneficial at higher throughput ranges. The 2740 IOA (which is targeted for smaller systems) has similar performance characteristics to the 2726 IOA over typical operating ranges, but has slightly slower performance at higher throughput ranges. These IOP/IOAs are also capable of attaching Ultra-SCSI (40 MB/sec bus) versions of the 6606, 6607, 6713 and 6714 DASDs. These devices can improve performance for workloads characterized by large disk I/O operations. The new 2748 PCI IOA has better performance characteristics, a larger write cache (26MB), and is capable of supporting the SCSI Wide-Ultra2 (80 MB/sec) bus.

The DASD that are used in the Internal DASD Subsystems have read ahead buffers that can provide performance advantages. Like the 9337, each of these DASD has a 512K buffer. The buffer is allocated into multiple segments that are larger than 32K each. Read ahead data from recent IOs are kept in these buffer segments. Depending on the data access patterns, it is possible that the data needed is already contained in a buffer segment. If so, no physical access to the DASD is required. Depending on your data access patterns, this can significantly improve performance. Our analysis of several specific customer installations indicates that 10% to 30% of their DASD IO for "interactive" transactions would have already been contained in the read ahead buffer. For "batch" type jobs, 25% to 45% of their DASD IO would have already been contained in the read ahead buffer. The RAMP-C workload being used in this section has less than 10% of it's DASD IOs already in the read ahead buffer.

For the 9337-2xx, 9337-4xx, 9337-5xx, 65x2(also 2726, 2740, 2741, 2748, 6533, 6751 and 6754) models in HA mode, most batch jobs will run nearly as fast as if they were run in "base" mode or mirrored mode. Only in extreme cases will the HA mode cause degradation. An example of the extreme case is when there are sequences of hundreds of writes to a single 9337 or 65x2 in a short period of time.

You must ensure that you have enough arms to support the volume of DASD IOs that your customer will require. In some situations, using the larger capacity DASD may result in an insufficient number of arms to handle the required DASD IO volume. The Capacity Planning tools should be used to verify your configuration.

The recommended threshold for maximum DASD utilization for 1 arm configurations is higher than the threshold for multiple arm configurations. The reason for the lower recommendation for multiple arms is that it is assumed that when 2 or more arms have an average utilization of 40%, some of the arms may be at the 50% - 55% range while others will be lower. QSIZE400 and BEST/1 allow a 1 arm configuration to reach 55% before they recommend that an additional DASD be added.

Consider the following example. Assume you are configuring a system and need approximately 4000 MB of DASD space. You have the choice of 4 x 988MB or 2 x 1967MB. The 4 x 988MB configuration will support approximately 70% more DASD IOs as the 2 x 1967MB configuration. Because there is a maximum number of DASD devices that can be attached to each model, using the larger drives will allow more MB of DASD to be configured on your system.

The Performance Monitor (STRPFRMON command) captures additional performance data (buffer hits, etc.) for the 65x2, 6533 and 6530 attached DASD. This data is available in the QAPMDISK performance data file and is documented in Appendix A of the *AS/400 Work Management V4R3* (SC41-5306-02).

14.2 DASD Performance - Interactive

The implementation of the 4KB page size on RISC will improve system DASD IO efficiency. As a result of the larger page size, some DASD subsystem interactive Ops/Sec/GB ranges will appear lower than on IMPI.

Some DASD system performance charts included for RISC may differ from similar charts published for IMPI. These performance differences can be attributed primarily to the following:

- Differences in system processor power
- Differences in main storage configurations
- Differences in system page size
- Differences in allocation of data and programs on DASD.

Therefore, direct comparisons between RISC and IMPI system DASD performance charts are not recommended.

DASD Subsystem Performance - Base or Mirrored

The following bar graphs compare the service times for the AS/400 DASD subsystem offerings. The IO operations being performed are 7KB transfer size, 70% are reads and 30% are writes, and 80% require a seek over 1/3 of the disk surface while 20% require no seek. Queueing time is not included.



Figure 14.1. DASD Subsystem Performance / Non-Raid Capable - Base Mode



Figure 14.2. DASD Subsystem Performance / Raid Capable - Base Mode

Conclusions / Recommendations

- The performance of 6606, 6607, 6713 and 6714 disks is faster with the 6533 IOP than with the 6512 IOP. The 6754, 6751, 6532, 2726, 2740 and 2741 DASD IOP/IOAs have performance characteristics similar to the 6533 IOP over typical operating ranges.
- The performance with 6606, 6607, 6713 and 6714 disks is faster than the previous DASD types for all the subsystems.
- The 9728 subsystem performance is faster than the 6530 subsystem for the same type of DASD.
- The 6512 subsystem performance is better than the 6502 subsystem for the same type of DASD. This is due primarily to a faster processor and a larger 4 MB Write cache in the 6512.
- The 6512 subsystem performance is slightly better than the 9337-5xx subsystem for the same type of DASD.
- The 6502 subsystem performance is significantly better (32%) than the 6530 subsystem for the same type of DASD. This is due primarily to the 2 MB Write cache in the 6502.
- "RAID Capable" DASD subsystems are faster in base mode than "Non-RAID Capable" due to their write cache.

- The potential effect of read-ahead buffers are shown for the cases of having 25% and 50% of the total disk operations already in the read ahead buffer. Depending on the data access patterns, the buffers may provide significant performance improvements.
- The above conclusions hold for batch environments also. For actual batch performance results refer to Table 14.1, Table 14.2, and Table 14.3.

AS/400 System Interactive Performance - Base

The following graph compares the relative interactive performance of an AS/400 model 510/2144 configured with 33.5GB of internal or external DASD. The internal load source drive was ignored for this comparison chart. The curves characterize what may occur on either a 'Base' configuration or a mirrored configuration. The graph compares the 9337-580 model with the 65x2-4GB models for a commercial interactive environment.



Figure 14.3. System Interactive Performance - Base Mode

Conclusions / Recommendations

- The 6533/6607 DASD provides better interactive performance than the 6512/6607 DASD.The 6754, 6751, 6532, 2726, 2740 and 2741 DASD IOP/IOAs have performance characteristics similar to the 6533 IOP over typical operating ranges.
- The 6512/6607 DASD provides better interactive performance than the 9337-580 DASD.

- The 6512/6607 DASD provides better interactive performance than the 6502/6607 DASD, especially at higher system throughput.
- Performance with the 6502/6607 DASD is slightly better than the 9337-580 DASD for lower throughput, but becomes worse at higher throughputs.
- The 65x2 and 9337-5xx configurations have reduced volumes of physical disk IO due to the write cache. The write cache also greatly improves the service time for write ops.
- The 65x2 write cache provides a significant performance advantage over the 6530. When a write is requested to a 65x2, the 65x2 writes the data to the write cache and to the nonvolatile cache backup and the application is allowed to continue. Through a combination of the write cache and 65x2 nonvolatile memory, the 65x2 ensures the integrity of the data even if a failure should occur.
- This graph is based on RAMP-C workload. Other environments may vary significantly.
- The RAMP-C benchmark's data access patterns are intentionally random, therefore, the read-ahead buffers provided only minimal benefit for RAMP-C. Depending on your data access patterns, the DASD read ahead buffers may provide significant performance improvements.
- Similar results may occur on other AS/400 models. Response time / throughput curves encounter a "knee" when a resource is used too heavily. CPU, main memory, IOP Processor and DASD are examples of resources that can cause "knees". If faster AS/400 CPUs are used, and other resources are unchanged, the possibility that memory or DASD will constrain the throughput increases. The BEST-1 Capacity Planner should be used to determine appropriate configurations.

AS/400 System Interactive Performance - Mirrored versus Base

The following graph compares the relative interactive performance of an AS/400 model 510/2144 configured with 8.4 GB of User DASD. The graph compares a mirrored environment with 16 arms to a base (not mirrored) environment with 8 arms. It also shows the system performance effects during the resync of a single arm.



Figure 14.4. System Interactive Performance - Mirrored versus Base - Internal DASD

Conclusions / Recommendations

- The mirrored configuration provides equal or better interactive performance than the base configuration (not mirrored). The better mirrored performance is due to having more arms to handle the larger number of read ops at higher throughput.
- The system performance is less during the time it takes to resync an arm, especially at higher throughput. The customer could choose to schedule the resync during a period of lower system activity or quiesce some applications during the resync time (20 to 40 minutes for a 1GB device). Larger devices will have proportionally longer resync time.
- This graph is based on RAMP-C workload. Other environments may vary significantly.

AS/400 System Interactive Performance - RAID

The following graph compares the 65x2 RAID DASD Subsystems with the 9337-580 HA Subsystem. All subsystems contained 8 4GB arms.



Figure 14.5. System Interactive Performance - RAID Mode

Conclusions / Recommendations

- The 6533/6607 RAID DASD provides better interactive performance than the 6512/6607 RAID DASD.The 6754, 6751, 6532, 2726, 2740 and 2741 DASD IOP/IOAs have performance characteristics similar to the 6533 IOP over typical operating ranges.
- Performance with the 9337-580 HA is comparable with the 6512/6607 RAID DASD for similar throughput.
- The 9337 measurements were done with 4 parity arms per array and the 65x2 measurements were done with 8 parity arms per array. In general 8 parity arms per array will provide better performance at higher throughputs. At low to medium throughput, there is little performance difference between 4 and 8 parity arms per an 8 arm array. On the 65x2(also 2726, 2740, 2741, 6533, 6751 and 6754), parity arrays of 8 or more arms should be configured with 8 parity arms if possible.

Impact of failed DASD in RAID Subsystem

This is a general discussion of RAID-5.

• 65x2(also 2726, 2740, 2741, 6533, 6751 and 6754) RAID and 9337 HA DASD subsystems let the AS/400 system continue to operate even after a single DASD failure.

- With system checksum, a DASD failure will cause the AS/400 system to stop and an IPL will be required.
- RAID-5 overhead can become significant when a DASD fails.

READ

- To read from a failed DASD, RAID-5 must read ALL remaining arms in the set. (This means anywhere from 3 to 9 overlapped reads, where 1 was sufficient before). This will have a significant effect on the **failed** DASD subsystem throughput and response time. This degraded mode will last until the DASD is repaired and the "rebuild" of the failed DASD's parity stripes are complete.
- Reads to other DASD on the same subsystem are unaffected.
- Bottom line, if the parity array has 4 arms, this results in 1.5 times increase in DASD IO read volume to this array. If the array has 8 arms, the result is 1.75 times increase in DASD IO read volumes to this array.

WRITE

- ◆ There are 3 separate scenarios that apply to RAID-5 writes with one failed DASD.
 - If the failed DASD is not involved (either for data or for the checksum stripe), the writes are handled as normal RAID-5 writes. (2 reads plus 2 writes)
 - With a write to a failed DASD, all remaining DASD in the set must be read and then one write will be done to the checksum stripe. (N-1 reads plus 1 write, N=number of DASD arms)
 - If the DASD that contains the checksum stripe is the failed DASD, then all that is required is a write to the DASD that contains your data. (1 write)
 - Bottom line, if the parity array has 4 arms, each write averages a 3.25 increase in DASD IO write volume to this array. If the array has 8 arms, the result is a 4.13 times increase in DASD IO write volumes to this array.

General discussion

- When running in "exposed" mode, the fewer the number of arms in each parity array, the smaller the degradation.
- On systems with smaller amounts of DASD capacity, the degradation will be more noticeable. This is because there are fewer arrays which means that a larger percentage of the DASD operations will be directed to the "exposed" array.
- The DASD IO to any subsystems that do not have a failed DASD are unaffected.
- If the Customer cannot tolerate the temporary performance degradation that would occur with a RAID-5 DASD failure, they should consider mirroring.
- To obtain acceptable performance with a failed RAID-5 DASD, some customers may have to delay nonessential work until after the DASD is repaired. For example, a customer may continue to process their on-line order entries but delay their office tasks.
- In configurations with small amounts of total DASD space and with high availability requirements, mirroring may be a more satisfactory option
- The estimated time to rebuild a DASD is approximately 30 minutes for a 8 arm array on a dedicated system with no other jobs running. If other concurrent jobs being run on the system are requesting 130 IOs per second to this DASD subsystem, the rebuild time will increase to approximately 1 hour.

The following chart compares the impact of one failed DASD on several configurations.



Figure 14.6. Performance Impact of Failed DASD

Conclusions / Recommendations

- The 6533 performs better in "exposed" mode than the 9337-580 due to the ability of the 6533 to handle higher throughput more efficiently. The 6754, 6751, 6532, 2726, 2740 and 2741 DASD IOP/IOAs have performance characteristics similar to the 6533 IOP over typical operating ranges.
- With a parity array of 8 arms on a 6533, "exposed" mode throughput is about half of normal throughput.
- If there were "n" 65x2s in the configuration, 1/n of the DASD IOs are to the "exposed" 65x2.
- As "n" gets larger, the impact of a DASD failure to overall system performance is reduced.
- Additional degradations occur during rebuild for both 65x2 and 9337. The rebuild can be scheduled for periods of lower system utilization.

Ops/Sec/GB Guidelines for DASD Subsystems

The metric used in determining DASD subsystem performance requirements is the number of I/O operations per second per installed GB of DASD (Ops/Sec/GB). Ops/Sec/GB is a measurement of throughput per actuator. Since DASD devices have different capacities per actuator, Ops/Sec/GB is used to normalize throughput for different capacities. An Ops/Sec/GB range has been established for each

DASD type so that if the DASD subsystem performance is within the established range, the average arm percent busy will meet the guideline of not exceeding 40%.

The implementation of the 4KB page size on RISC will improve system DASD IO efficiency. As a result of the larger page size, some DASD subsystem interactive Ops/Sec/GB ranges will appear lower than IMPI.

The following bar charts show the "rule of thumb" for the Physical system Ops/Sec/GB of usable space that internal DASD subsystems can achieve with various DASD types. (To compute usable GB, we assume that the DASD subsystems have 8 disk units installed). The top of each bar is the volume of 7K transfer, 80% 1/3 seek, 30% write operations that each model can achieve when it is 40% busy. For the 6502, we assume that the write cache has an efficiency of 55%. For the 6512, 6532,6533, 6751, 6754, 2726, 2740 and 2741 we assume that the write cache has an efficiency of 65%. The vertical scale is the volume of physical Ops issued from the system. The 65x2 RAID disk subsystem bars are lower because of the additional work that these subsystems must do to maintain the RAID-5 parity stripes.



Figure 14.7. Ops/Sec/GB - Internal DASD



Figure 14.8. Ops/Sec/GB - Internal DASD (continued)

Conclusions / Recommendations

- In general, the higher the capacity of the DASD device the lower it's throughput range will be. The 6714 device (17548 MB per arm) has a lower throughput range than the 6606 device (1967 MB per arm).
- Typically, DASD subsystems will have a lower throughput range when operated in RAID mode rather than Base mode. The throughput difference due to RAID will tend to be smaller for workloads characterized by higher read to write ratios.
- The 6714 DASD (17548 MB per arm) is more appropriate when the capacity requirement is very large and the Ops/Sec/GB requirement is less than 3.
- The 6713 DASD (8589 MB per arm) is more appropriate when the capacity requirement is large and the Ops/Sec/GB requirement is less than 4.
- The 6607 models (4194 MB per arm) will be the appropriate choice for almost all other situations. *NOTE: The 6606 models (1976 MB per arm) may be needed for extreme cases where the capacity requirement is very low relative to the system disk ops/sec rate. Refer to section OSGB for a discussion of the performance limits for each DASD type. 6606 is only needed for cases exceeding the limits for 6607.

- The low cost 9728 DASD IOA provides a similar throughput range when compared to the 6530 DASD IOP.
- The 6512 DASD subsystem provides an improved throughput operating range when compared with the 6502 DASD subsystem.



Figure 14.9. Ops/Sec/GB - Internal versus External DASD

Conclusions / Recommendations

- The 6533 DASD subsystem provides an improved throughput operating range when compared with the 6512 DASD subsystem. The 6754, 6751, 6532, 2726, 2740 and 2741 DASD IOP/IOAs have performance characteristics similar to the 6533 IOP over typical operating ranges.
- The 6533 DASD subsystem provides an improved throughput operating range when compared with the 9337-5xx DASD subsystem.
- The 6533 DASD IOP provides a better throughput range when compared to the low cost 9728 DASD IOA.

Using the OPS/SEC/GB Chart

The Ops/Sec/GB chart above should be used as a guideline on what DASD model is the appropriate choice when adding or upgrading DASD. In conjunction with this chart, you should utilize results obtained from

the Performance Tools LPP report to determine what DASD model meets your DASD performance requirements. For more detailed DASD performance analysis, it is recommended to use BEST/1-400, the capacity planner for the AS/400.

Operations per second is a measurement of throughput per actuator. Since DASD devices have different capacities per actuator, operations per second per GB is used to normalize throughput for different capacities. To determine the operations per second of your current operating environment, follow the procedure outlined below. The value obtained by this procedure will help determine what DASD model will meet your current or projected DASD performance requirements.

- 1. Collect performance data using the Performance Monitor. Be sure to collect this data during peak activity for at least a one hour time period using 10 minute sample intervals.
- 2. Print the Performance Tools System Report using the PRTSYSRPT command. Then, refer to the "Disk Utilization" section of the Performance Tools LPP System Report. From this report the following data can be obtained:
 - Total operations per second Use Op Per Second column
 - Total GBs of DASD installed use Size (M) column
- 3. To determine the total GBs installed, simply add the "Size (M)" column and divide by 1000. When adding the total GBs, you should ONLY include the disk units you plan to replace. Also, if Mirroring is active, divide the total GB being mirrored by 2 when calculating the sum.
- 4. To determine the total operations per second, add the total operations per second number ("Op Per Second" column). When adding the total operations per second, you should ONLY include the disk units you plan to replace. Also, if mirroring is active, you need to divide the total number of operations per second for all mirrored units by 2.
- 5. To determine the operations per second per GB, divide the total operations per second you calculated in step 4, by the total GBs installed value you calculated in step 3.

You can then use the operations per second value to determine what model of DASD best fits your current or projected DASD performance requirement.

EXAMPLE

If you wanted RAID mode, you could select 6533/6607 models if the physical system IO/Sec/GB were at 11 or fewer.

14.3 DASD Performance - Batch

Commercial Batch - Base versus RAID

This section shows the results of running one of IBM's batch workloads in a dedicated environment. The workload performs the following functions :

- Sequential and Keyed Record Copy
- Sequential and Keyed Program Read
- Sequential and Keyed Record Read/Update
- Record Matching
- Adding and Removing Members
- RGZPFM of 500,000 Records
- Average 40% Read Ops, 60% Write Ops, 17 KB/IO
- 70% Synchronous / 30% Asynchronous Ops

The workload was run in a 24 MB memory pool. A 6533-18GB was compared to an equivalently configured 6532-18GB with RAID protection turned on and off.



Figure 14.10. Commercial Batch Performance - Base versus RAID

Conclusions / Recommendations

- The 6533-18GB (6714) has better performance than the 6532-18GB (6614) for this commercial batch workload.
- The 6714 DASD has slightly better performance than the 6713 DASD.



The workload was run in a 24 MB memory pool. A 6512-8GB was compared to an equivalently configured 9337-590 with RAID protection turned on and off.

Figure 14.11. Commercial Batch Performance - Base versus RAID

Conclusions / Recommendations

- The 6512-8GB (6613) has better performance than the 6502-8GB (6613) for this commercial batch workload.
- The 6512-8GB (6613) has better performance than the 9337-590.
- The 6502-4GB (6607) has similar performance as the 9337-580 for this commercial batch workload.
- The 6607 DASD has slightly better performance than the 6713 DASD.

14.4 DASD Performance - General

Mixing RAID DASD with other DASD in one ASP

Combining 65x2(also 2726, 2740, 2741, 6533, 6751 and 6754) RAID DASD with mirrored DASD in a single ASP is allowed. Combining RAID DASD with mirrored DASD on the same 65x2 is also allowed.

Write Intensive Applications (eg RESTORE)

RAID-5 (like system checksum) can have a significant impact on batch type programs that issue many writes in a short period of time. This is due to the four times increase in disk IO required for each write. The 65x2 write cache handles this impact for almost all scenarios except those that write hundreds of writes to the DASD in a very short period of time. Even in this worst case scenario, with only one 65x2 array configured, the restore of a large file took only 30% longer than a restore to a standard" 6530 (or 9728) configuration.

The 65x2 can restore small objects faster than 6530 because of the 65x2 write cache. The write cache provides fast completion of write requests and is able to "stay ahead" of the system.

The 65x2 RAID models offer significant advantages in availability, reliability, price, etc. One of the "costs" of the availability advantage is the increased time to restore data. This increase in time needs to be considered when planning the installation of RAID-5 disk. The time to load data onto the RAID-5 boxes must be included in the overall installation planning. With the increased availability and reliability offered by 65x2 RAID, the necessity to reload the data again due to a single disk unit failure will be eliminated.

DST "Add Unit"

Part of the process of adding DASD to a system is using the Dedicated Service Tool (DST) to "Add Unit". This ensures that the entire DASD(s) is initialized with a X'00' data pattern and verified.

When multiple DASD are added at once, the system will add up to 16 units in parallel.

The time for adding up to 16 units on a DASD IOP (Base mode) is approximately

- 48 minutes for 2 GB arms (6606)
- 86 minutes for 4 GB arms (6607)
- 162 minutes for 9 GB arms (6713)
- 302 minutes for 18 GB arms (6714)

The Dedicated Service Tool is also used to start and stop parity (RAID-5) arrays on the 65x2 (also 2726, 2740, 2741, 6751 and 6754) IOP. When a parity array is initially set up, the fastest approach is to start parity on an array first and then add the arms to an ASP. The time required for this process (start parity and add) on two 8 arm arrays is approximately

- 48 minutes for 2 GB arms (6606)
- 86 minutes for 4 GB arms (6607)
- 162 minutes for 9 GB arms (6713)
- 302 minutes for 18 GB arms (6714)

If the arms are added to the ASP before starting the array, then the time required may double. If a system IPL occurs between starting the array and adding the arms to an ASP, then the time required could be 3 times as long.

If the arms are currently part of an ASP, then starting an array will take longer because the system may need to move data before it synchronizes the parity stripes. This could take up to

- 90 minutes for 2 GB arms (6606)
- 160 minutes for 4 GB arms (6607)
- 300 minutes for 9 GB arms (6713)
- 580 minutes for 18 GB arms (6714)

Stopping parity on an 8 arm array takes about

- 70 seconds for 2 GB arms (6606)
- 120 seconds for 4 GB arms (6607)
- 220 seconds for 9 GB arms (6713)
- 410 seconds for 18 GB arms (6714)

14.5 Integrated Hardware Disk Compression (IHDC)

Integrated Hardware Disk Compression (IHDC) is a new DASD capability for V4R3. IHDC has the following characteristics :

- Data is dynamically compressed/decompressed by the DASD subsystem controller (IOP/IOA) independent of the AS/400 system processor
- Compressed data is not seen above the DASD controller level
- Compression is performed by an LZ1 compression chip on the DASD controller
- An average 2X compression ratio, with up to 4X achievable (data dependent)
- Customer on/off option provided at disk arm level
- RAID and mirroring is supported (no additional restrictions)
- With compression, AS/400 disk capacity maximums can be exceeded.

IHDC provides the following customer value :

- Reduces cost of on-line storage
- Provides better performance than typical software data compression
- Enables new applications
- Protects investment ability to increase capacity of installed DASD
- Provides a storage management solution when used with hierarchical storage management (HSM).

IHDC has the following requirements :

- Requires a compression enabled IOP/IOA
- Compressed DASD must be configured in user ASPs only
- 17.54GB disks are supported with V4R4 and future releases
- Disks must be unconfigured to be enabled for compression
- Data must be saved before compression is disabled to avoid loss
- Compressed disks can be migrated only to compression enabled IOP/IOAs.

Disk Compression Positioning:											
High	Disk Compression Recommended	Disk Compression Candidate									
Compressibility	Disk Compression Candidate	High Performance Disk Recommended									
	Low Throu	ughput High									

- High compressibility, low throughput
 - The disk space reduction benefit of compression will be realized, and both reads and writes to compressed disk should outperform uncompressed disk
 - Disk compression is recommended
- Low compressibility, low throughput
 - The lower the compressibility, the smaller the disk space reduction benefit, although disk performance will likely not suffer
 - Disk compression is a candidate if compressibility is high enough
- High compressibility, high throughput
 - The disk space reduction benefit of compression will be realized, and compressed disk performance may be close to uncompressed disk performance, or even faster for read intensive applications
 - ♦ Disk compression is a candidate, especially for read intensive applications
- Low compressibility, high throughput

- The disk space reduction benefit of compression will be minimal, and performance will likely suffer especially for write intensive applications
- ✤ High performance disk (uncompressed) is recommended

DASD Compression Performance Guidelines

Precise performance projections for IHDC are not possible due to :

- Compressibility of data which can vary greatly
- Workload characteristics of specific applications

DASD compression may cause performance to vary, in general :

- System performance impacts will be minimal when DASD operations are light
- For data with low compression rates (< 2X), DASD read/write performance will generally be slower than for uncompressed data
- For data with high compression rates (> 3X), DASD read/write performance can be faster than for uncompressed data
- DASD Read intensive workloads will typically perform better than DASD Write intensive workloads
- Interactive applications with a mixture of DASD reads and writes with medium to heavy DASD operations should use high performance uncompressed DASD
- With compressed disks it is critical that they operate within a reasonable margin below 'dasd full', otherwise performance will be greatly affected. In contrast to uncompressed disks, when a compressed disk approaches full (approximately 85%) a disk defragmentation task is started within the IOP/IOA to recover fragmented storage and may take considerable time to finish. Since this task runs concurrently with system operations, performance will be degraded until this task completes.
- Mixing Compressed DASD with Uncompressed DASD on the same IOP/IOA may impact the performance of the Uncompressed DASD due to higher IOP/IOA utilization.
- Mixing Compressed DASD with Uncompressed DASD within the same user ASP is supported but is not recommended due to potential performance impact caused by unbalanced disk utilization.

DASD compression is intended for :

- Vast amounts of historical or archive data
- Low activity data
- Spool files

- Journals
- Save files (staging)

Not for highly volatile data or data already compressed (images, etc.)

Types of applications that can benefit from DASD compression :

- Data warehouse, data mining
- On-line access to archive data
- On-line viewing of reports (paper or micro-fiche replacement)
- Part of a hierarchical storage management strategy

Applications are a candidate for DASD compression if :

- Additional DASD storage is required
- Application data can be partitioned (at least partially) into user ASPs
- Top application performance is not required

Refer to AS/400 Backup and Recovery V4R3 (SC41-5304-02) for more information about configuring and using compressed DASD.

Interactive Performance with Compressed DASD

The following graph compares relative interactive system performance of an AS/400 model 640/2239 configured with 16-arm user ASPs of Compressed, Uncompressed, RAID/Compressed and RAID/Uncompressed DASD. The graph compares the performance results of running the RAMP-C workload in each of the 4 user ASPs.



Figure 14.13. System Interactive Performance - Compressed DASD

- At lower throughputs and with equal number of disk arms, Compressed DASD has similar (0-10% degradation) system performance characteristics to Uncompressed DASD for interactive workloads.
- Compressed DASD is not appropriate for high throughput (ie. write intensive) environments.
- Compressed DASD performance can often be improved by maintaining :
 - Disk CPU utilization below 60% (about 360 ops/sec per IOP/IOA)
 - Disk utilization below 40% (about 23 ops/sec per disk arm)
- For the above graph, the Compressed disk 40% utilization point is at 146,000 transactions per hour and the RAID/Compressed disk 40% utilization point is at 140,000 transactions per hour.

- For the above graph, the Compressed Disk CPU (IOP/IOA) 60% utilization point is at 144,000 transactions per hour and the RAID/Compressed Disk CPU 60% utilization point is at 138,000 transactions per hour. Above these limits, system performance with Compressed DASD tends to degrade noticeably as the throughput increases.
- Compressed DASD with RAID has similar system performance characteristics to Compressed DASD without RAID at lower throughputs. At higher throughputs, RAID/Compressed DASD performance is less due to higher utilization's for the same op rates. The same criteria as above should be followed for obtaining acceptable RAID/Compressed DASD performance.
- Just as with uncompressed DASD, the number of disk arms must be adequate to support anticipated op rates.
- Configuring fewer disk arms per IOP/IOA will typically improve the performance of Compressed DASD. DASD subsystems with 8 arms per IOP/IOA will usually perform much better than those with 16 arms per IOP/IOA.

Batch Performance with Compressed DASD

The following chart compares system performance of various batch type applications while running on an AS/400 model S30/2259 configured with 16-arm user ASPs of Compressed, Uncompressed, RAID/Compressed and RAID/Uncompressed DASD. Batch run time was measured in each of the 4 user ASPs for 7 batch tests with the following DASD I/O characteristics :

- 1. Sequential read ops, 5 KB/op, OS/400 Expert Cache off
- 2. Sequential read ops, 60 KB/op, OS/400 Expert Cache on
- 3. Sequential read and write ops, 68% reads, 5 KB/op, OS/400 Expert Cache off
- 4. Sequential read and write ops, 17% reads, 50 KB/read op, 5 KB/write op, OS/400 Expert Cache on
- 5. Random read ops, 7 KB/op, OS/400 Expert Cache off
- 6. Random write ops, 8 KB/op, OS/400 Expert Cache off
- 7. Sequential read and write ops, 14% reads, 5 KB/op, OS/400 Expert Cache off



Figure 14.14. Batch Run Time Performance - Compressed DASD

- For batch applications characterized by DASD read ops, system performance varied only slightly between Compressed, Uncompressed, RAID/Compressed and RAID/Uncompressed DASD ASPs.
- For batch applications characterized by DASD write ops, system performance was slower for Compressed and RAID/Compressed than for Uncompressed and RAID/Uncompressed DASD ASPs.
- For batch applications characterized by a mixture of DASD read and write ops, system performance was slower for Compressed and RAID/Compressed than for Uncompressed and RAID/Uncompressed DASD ASPs. The magnitude of the performance difference typically depends on the percentage of write ops.
- OS/400 Expert Cache provided better batch system performance when active.

Save/Restore Performance with Compressed DASD

The following charts compare system performance of Save/Restore operations while running on an AS/400 model S30/2259 configured with 16-arm user ASPs of Compressed, Uncompressed, RAID/Compressed and RAID/Uncompressed DASD. The System ASP (ASP1) was configured with Uncompressed DASD. Data transfer rates were measured in each of the 4 user ASPs for 6 different types of Save/Restore tests :

- 1. SAV to a 3590 tape from the user ASP (Read data from Compressed DASD)
- 2. RST from a 3590 tape to the user ASP (Write data to Compressed DASD)
- 3. SAV from ASP1 to a *SAVF on the user ASP (Write data to Compressed DASD)
- 4. RST from a *SAVF on the user ASP to ASP1 (Read data from Compressed DASD)
- 5. SAV to a *SAVF on ASP1 from the user ASP (Read data from Compressed DASD)
- 6. RST from a *SAVF on ASP1 to the user ASP (Write data to Compressed DASD)

The data used for the first chart has a compression ratio of 2X and the data for the second chart has a compression ratio of 4X.



Figure 14.15. Compressed DASD Save/Restore Rates - 2X Compressibility



Figure 14.16. Compressed DASD Save/Restore Rates - 4X Compressibility

- For Save operations to tape, Compressed DASD (also RAID/Compressed DASD) has approximately the same system performance characteristics as Uncompressed DASD. Save operations primarily issue read op to DASD and read op performance is very similar for Compressed, RAID/Compressed and Uncompressed DASD.
- For Restore operations from tape, Compressed DASD performance is highly dependent upon the compressibility of the data being restored the better the data compresses the better the restore performance.
 - Performance can range from 50% degradation (2x compression) to 10% degradation (4x compression) for Compressed DASD compared to Uncompressed DASD.
 - Performance can range from 50% degradation (2x compression) to 40% improvement (4x compression) for RAID/Compressed DASD compared to RAID DASD. The improvement is due to better write cache efficiency and smaller ops because of data compression.
 - Restore operations primarily issue large write ops to DASD along with allocate ops. Performance degradation occurs because ops have higher overhead due to compression and generate higher utilization for the IOP/IOA and devices.

Data Migration Performance with Compressed DASD

The following chart compare system performance of BRMS data migration operations while running on an AS/400 model S30/2259 configured with 16-arm user ASPs of Compressed and Uncompressed DASD. The System ASP (ASP1) was configured with Uncompressed DASD. Data transfer rates were measured between ASP1 and each of the 2 user ASPs for 4 different types of data :

- Large data base file
- Typical user mix of data
- Source data
- DLO data



Figure 14.17. Compressed DASD Data Migration Rates

Conclusions / Recommendations

• Migrating data from an Uncompressed ASP to a Compressed ASP is usually slower than migrating data between two Uncompressed ASPs. The magnitude of the performance difference typically depends on the type and compressibility of the data being migrated, the lower the compressibility, the slower the transfer rate.

• Migrating data from a Compressed ASP to an Uncompressed ASP is usually as fast as or slightly faster than migrating data between two Uncompressed ASPs.

Data Throughput Performance with Compressed DASD

The following graph compares relative data throughput performance of Compressed and Uncompressed DASD read and write operations for the range of data compression ratios supported. This graph is intended to show in general how data throughput for Compressed DASD is dependent upon the compressibility of the data being transferred.



Figure 14.18. Compressed DASD Relative Data Throughput

- Performance of Uncompressed DASD read and write operations is independent of the compression ratio of the data.
- As compression ratios improve, performance improves accordingly, with read always leading write performance .
- At high compression ratios, compression performance can actually exceed uncompressed performance.

Response Time Performance with Compressed DASD

The following graph compares relative response time performance of Compressed and Uncompressed DASD read and write operations for the range of data throughput. This graph is intended to show in general how response time for Compressed DASD depends upon the data throughput of the DASD devices.



Figure 14.19. Compressed DASD Relative Response Time

- As data throughput increases, DASD response time increases until it reaches a bottleneck at higher throughputs.
- Compression DASD tends to bottleneck sooner than Uncompressed DASD.
- Writes tend to bottleneck sooner than Reads.

14.6 DASD Subsystem Performance Improvements for V4R4

This section discusses the DASD subsystem performance improvements that are new for the V4R4 release. These consist of the following new hardware and software offerings :

- PCI RAID Disk Unit Controller (#2748)
- 10K RPM Disks (#6717)
- Storage/PCI Expansion Tower (#5065)
- Extended Adaptive Cache

The PCI RAID Disk Unit Controller (#2748) is a new DASD IOA that attaches to the system PCI bus. It provides performance improvements by increasing Fast Write Cache to 26 MB (from 4MB) and adding SCSI LVD (Low Voltage Differential Signaling) for SCSI Wide-Ultra2 (80MB) support on a new storage adapter. When the 2748 IOA is configured for DASD Compression the Fast Write Cache is limited to 4 MB.

The 6717 is a new 10K RPM Disk (9GB) that provides faster data access that the previous 7200 RPM devices. It can be attached only with 6532, 6533, 6751, 6754, 2726, 2740, 2741 and 2748 IOP/IOAs. It can be used as a load source and can be RAIDED and MIRRORED with its 7200 RPM counterparts.

The Storage/PCI Expansion Tower (#5065) provides connectability of the new PCI RAID Disk Unit Controller and 10K RPM Disks to a system which has SPD buses only.

The Extended Adaptive Cache is a feature of the new PCI RAID Disk Unit Controller that provides improved performance characteristics especially in read intensive workloads. The Extended Adaptive Cache requires a Read Cache Device (#4331 or #6831) for memory. For V4R4, the Read Cache Device is a 1.6GB volatile solid state disk. The Extended Adaptive Cache is managed such that ranges of data actively being read are brought into and kept in the cache for as long as they remain active. The goal is to improve performance for read-only or read-write commercial type workloads, while not harming the performance of write-only, random, sequential read, or sequential write workloads.

Extended Adaptive Cache was created to complement other caches within the system and designed to meet the specific needs of AS/400 system users. Although Extended Adaptive Cache functions independently from Expert Cache (which uses main memory), the DASD IOA fast write cache, and device read-ahead buffers, it takes each caching strategy into account as it tracks the physical I/O activity. NOTE: DASD Compression and Extended Adaptive Cache are mutually exclusive.

Although Extended Adaptive Cache has proven to be highly effective in improving performance on many types of workloads, the cache effectiveness is workload dependent. Both the system configuration and type of I/O activity have a direct impact on the performance benefits of Extended Adaptive Cache. Therefore the Extended Adaptive Cache Simulator was created to enable AS/400 system users to know what benefits would be realized by adding Extended Adaptive Cache to their system, before having to purchase the actual cache memory. Extended Adaptive Cache Simulator is an integrated performance evaluation tool that uses the same algorithms that would manage Extended Adaptive Cache if it were activated. Refer to the Extended Adaptive Cache white paper at http://www.as400.ibm.com/hsmcomp/EACacheWhitepaper.htm for additional information.

DASD Subsystem Performance - PCI RAID Controller (#2748) and 10K RPM Disks

The following bar graph compares the service times for the new AS/400 DASD subsystem offerings. The new 2748 IOA is compared to the previous 2741 IOA and the new 10K RPM 6717 disk is compared the previous 7200 RPM 6607, 6713 and 6714 disks. The IO operations being performed are 7KB transfer size, 70% are reads and 30% are writes, and 80% require a seek over 1/3 of the disk surface while 20% require no seek. Queuing time is not included.



Figure 14.20. DASD Subsystem Performance - PCI RAID Controller (#2748) and 10K RPM Disks

- The performance of previous 6607, 6713 and 6714 disks is faster with the new 2748 IOA than with the previous 2741 IOA. This performance improvement can be up to 20% faster in some cases.
- The performance on the 2741 IOA with the new 10K RPM 6717 disk is faster than with the previous 7200 RPM 6607, 6713 and 6714 disks. This performance improvement can be up to 20% faster in some cases.

- The performance on the new 2748 IOA with the new 10K RPM 6717 disk is faster than with the previous 7200 RPM 6607, 6713 and 6714 disks on the previous 2741 IOA. This performance improvement can be up to 60% faster in some cases.
- The potential effect of device read-ahead buffers are shown for the cases of having 25% and 50% of the total disk operations already in the read ahead buffer. Depending on the data access patterns, the buffers may provide significant performance improvements.
- The above conclusions hold for batch environments also. For actual batch performance results refer to Table 14.2.

AS/400 System Interactive Performance - PCI RAID Disk Unit Controller (#2748)

The following graph compares the relative interactive performance of an AS/400 model 720/2064(1505) configured with 14 10K RPM (6717) DASD when running the CPW workload. The internal load source drive was ignored for this comparison chart. The curves characterize what may occur on either a 'No Protection' (also Mirrored) or RAID configuration. The graph compares the new 2748 IOA with the previous 2741 IOA.



Figure 14.21. System Interactive Performance - PCI RAID Disk Unit Controller (#2748)

- The new 2748 PCI DASD IOA provides significantly better interactive performance than the previous 2741 DASD IOA.
- When operating with no DASD protection and at 40% DASD utilization, the system throughput improves by approximately 27% and the DASD subsystem throughput improves by about 45%.
- When operating with RAID protection and at 40% DASD utilization, the system throughput improves by approximately 28% and the DASD subsystem throughput improves by about 38%.
- The 2748 write cache (26MB) provides a significant performance advantage over the 2741 write cache (4MB).
- This graph is based on CPW workload. Other environments may vary significantly.
- The CPW benchmark's data access patterns are intentionally random, therefore, the read-ahead buffers provided only minimal benefit for CPW. Depending on your data access patterns, the DASD read ahead buffers may provide significant performance improvements.
- Similar results may occur on other AS/400 models. Response time / throughput curves encounter a "knee" when a resource is used too heavily. CPU, main memory, IOP Processor and DASD are examples of resources that can cause "knees". If faster AS/400 CPUs are used, and other resources are unchanged, the possibility that memory or DASD will constrain the throughput increases. The BEST-1 Capacity Planner should be used to determine appropriate configurations.

AS/400 System Interactive Performance - 10K RPM Disks

The following graph compares the relative interactive performance of an AS/400 model 720/2064(1505) configured with 14 DASD when running the CPW workload. The internal load source drive was ignored for this comparison chart. The curves characterize what may occur on a RAID configuration. The graph compares the new 10K RPM Disks with the previous 7200 RPM Disks.



Figure 14.22. System Interactive Performance - 10K RPM Disks

- The new 10K RPM Disks (6717) provides significantly better interactive performance than the previous 7200 RPM Disks (6713).
- When operating with RAID protection on the 2741 IOA and at 40% DASD utilization, the system throughput improves by approximately 20% and the DASD subsystem throughput improves by about 25%.
- This graph is based on CPW workload. Other environments may vary significantly.
- The CPW benchmark's data access patterns are intentionally random, therefore, the read-ahead buffers provided only minimal benefit for CPW. Depending on your data access patterns, the DASD read ahead buffers may provide significant performance improvements.
- Similar results may occur on other AS/400 models. Response time / throughput curves encounter a "knee" when a resource is used too heavily. CPU, main memory, IOP Processor and DASD are examples of resources that can cause "knees". If faster AS/400 CPUs are used, and other resources are

unchanged, the possibility that memory or DASD will constrain the throughput increases. The BEST-1 Capacity Planner should be used to determine appropriate configurations.

AS/400 System Interactive Performance - DASD Compression

The following graph compares the relative interactive performance of an AS/400 model 720/2064(1505) configured with the new 2748 IOA and 14 10K RPM (6717) DASD when running the CPW workload. The internal load source drive was ignored for this comparison chart. The curves characterize what may occur on this DASD subsystem configuration. The graph compares the DASD Compression/RAID with 'No Protection' and RAID Protection.



Figure 14.23. System Interactive Performance - DASD Compression

- At lower throughputs and with equal number of disk arms, RAID/Compressed DASD has similar (0-10% degradation) system performance characteristics to RAID/Uncompressed DASD for interactive workloads.
- When operating with RAID protection and at 40% DASD utilization, the system throughput with Compression turned on is approximately 14% less than with Compression turned off.
- With the new 2748 PCI DASD IOA and 10K RPM Disks, the RAID/Compressed DASD can operate at higher throughputs than with previous 2741 DASD IOA and 7200 RPM Disks before its

performance begins to deviate from RAID DASD. Just as with uncompressed DASD, the number of disk arms must be adequate to support anticipated op rates.

• This graph is based on CPW workload. Other environments may vary significantly.

AS/400 System Interactive Performance - Extended Adaptive Cache

The following set of graphs compare the relative interactive performance of an AS/400 model 720/2064(1505) configured with the new 2748 IOA and 14 10K RPM (6717) DASD when running the CPW workload. The internal load source drive was ignored for these comparison charts. The curves characterize what may occur on this RAID DASD subsystem configuration when the Extended Adaptive Cache is enabled with a 1.6GB Read Cache Device. The first graph compares the Extended Adaptive Cache on and off with the system level Expert Cache (EC) on and tuned for the CPW workload by using the *Usrdrn option. The third graph compares the Extended Adaptive Cache is not option. The third graph compares the Extended Adaptive Cache on and off with the system level Expert Cache (EC) workload.



Figure 14.24. System Interactive Performance - Extended Adaptive Cache without Expert Cache







Figure 14.26. System Interactive Performance - Extended Adaptive Cache with Expert Cache

- The Extended Adaptive Cache (EAC-On) provides improved interactive performance whether Expert Cache is turned on or not (EC-On or EC-Off). The Extended Adaptive Cache is automatically enabled for each 2748 IOA whenever a Read Cache Device (1.6 GB solid state disk) is attached to one of its SCSI busses. Expert Cache is activated in a system memory pool by issuing an OS/400 command.
- With Expert Cache off, Extended Adaptive Cache provides significantly better interactive performance. At lower throughputs, system response time is cut in half. When operating at 40% DASD utilization, the system throughput improves by approximately 17% and the DASD subsystem throughput improves by about 26%.
- With Expert Cache on (and tuned for the CPW workload by using the *Usrdfn option), Extended Adaptive Cache provided additional interactive performance improvements over and above that given by Expert Cache. At lower throughputs, system response time was only slightly better but got faster at higher throughputs. When operating at 40% DASD utilization, the system throughput improves by approximately 6% and the DASD subsystem throughput improves by about 9%.
- With Expert Cache on (and NOT tuned for the CPW workload), Extended Adaptive Cache provided interactive performance improvements similar to that with Expert Cache off. At lower throughputs, the system response time is much faster but at higher throughputs the difference is smaller. When operating at 40% DASD utilization, the system throughput improves by approximately 20% and the DASD subsystem throughput improves by about 24%.
- This graph is based on CPW workload. Other environments may vary significantly.

Ops/Sec/GB Guidelines for PCI RAID Controller (#2748) and 10K RPM Disks

The metric used in determining DASD subsystem performance requirements is the number of I/O operations per second per installed GB of DASD (Ops/Sec/GB). Ops/Sec/GB is a measurement of throughput per actuator. Since DASD devices have different capacities per actuator, Ops/Sec/GB is used to normalize throughput for different capacities. An Ops/Sec/GB range has been established for each DASD type so that if the DASD subsystem performance is within the established range, the average arm percent busy will meet the guideline of not exceeding 40%.

The following bar charts show the "rule of thumb" for the Physical system Ops/Sec/GB of usable space that internal DASD subsystems can achieve with various DASD types. (To compute usable GB, we assume that the DASD subsystems have 8 disk units installed). The top of each bar is the volume of 7K transfer, 80% 1/3 seek, 30% write operations that each model can achieve when it is 40% busy.



Figure 14.27. Ops/Sec/GB - PCI RAID Controller (#2748) and 10K RPM Disks

- The new 2748 DASD subsystem provides an improved throughput operating range when compared with the 2741 DASD subsystem for both 'Base' and RAID mode.
- The new 10K RPM 9GB Disk (6717) provides an better throughput operating range when compared with the 7200 RPM 9GB Disk (6713).

Batch Performance - PCI RAID Disk Unit Controller (#2748)

The following chart compares system performance of various batch type applications while running on an AS/400 model 720/2064(1505) configured with 14-arm user ASPs of RAID DASD. One ASP used a 2748 IOA and the other ASP used a 2741 IOA. The new 10k RPM 6717 disks were used for all of these measurements. Batch run time was measured for both of the user ASPs for 7 batch tests with the following DASD I/O characteristics :

- 1. Sequential read ops, 5 KB/op, OS/400 Expert Cache off
- 2. Sequential read ops, 60 KB/op, OS/400 Expert Cache on

- 3. Sequential read and write ops, 68% reads, 5 KB/op, OS/400 Expert Cache off
- 4. Sequential read and write ops, 17% reads, 50 KB/read op, 5 KB/write op, OS/400 Expert Cache on
- 5. Random read ops, 7 KB/op, OS/400 Expert Cache off
- 6. Random write ops, 8 KB/op, OS/400 Expert Cache off
- 7. Sequential read and write ops, 14% reads, 5 KB/op, OS/400 Expert Cache off



Figure 14.28. Batch Run Time Performance - PCI RAID Disk Unit Controller (#2748)

- The new 2748 PCI DASD IOA provides better system performance than the previous 2741 DASD IOA for all 7 of the batch tests. For tests that had write operations, the performance was significantly better. This can be attributed primarily the bigger fast write cache (26MB).
- OS/400 Expert Cache provided better batch system performance when active.

Batch Performance - DASD Compression

The following chart compares system performance of various batch type applications while running on an AS/400 model 720/2064(1505) configured with 14-arm user ASPs of Compressed, Uncompressed, RAID/Compressed and RAID/Uncompressed DASD. The new 2748 IOA and 10k RPM 6717 disks were used for these measurements. Batch run time was measured in each of the 4 user ASPs for the 7 batch tests.



Figure 14.29. Batch Run Time Performance - DASD Compression

- The new 2748 PCI DASD IOA and 10K RPM Disks provide better system performance for all 7 of the batch tests than was measured previously with 7200 RPM disks and the older DASD IOA (see header Batch Performance with Compressed DASD in section 14.5).
- For tests that had write operations, the performance was significantly better. This can be attributed primarily the bigger fast write cache (26MB).
- OS/400 Expert Cache provided better batch system performance when active.

Batch Performance - Extended Adaptive Cache

The following chart compares system performance of various batch type applications while running on an AS/400 model 720/2064(1505) configured with 14-arm user ASP of RAID DASD. One ASP had the Extended Adaptive Cache enabled and the other ASP did not. The new 2748 IOA and 10k RPM 6717 disks were used for these measurements. Batch run time was measured for both of the user ASPs for the 7 batch tests.



Figure 14.30. Batch Run Time Performance - Extended Adaptive Cache

- For batch applications characterized by sequential read ops or random read ops, system performance was either only slightly less or greater with Extended Adaptive Cache enabled. This can be attributed to the fact that Extended Adaptive Cache is designed to improve performance for read-only or read/write commercial type workloads while not harming the performance of write-only, random, sequential read, or sequential write workloads. The algorithms used in Extended Adaptive Cache are NOT read-ahead based, but instead rely on locality of reference, a characteristic not exhibited by these batch applications.
- OS/400 Expert Cache provides better performance than Extended Adaptive Cache when the applications are characterized by primarily sequential read ops and should be used in these situations. The algorithms used in Expert Cache can roughly be described as read-ahead algorithms.

• OS/400 Expert Cache provided better batch system performance when active.

Chapter 15. Save/Restore Performance

Many factors influence the observable performance differences of save operations and restore operations. These factors include:

- Hardware (such as tape drives)
- The tape input/output processor (IOP) that is used
- Placement of the tape IOP in the system
- Type of workload (Large file, User Mix)

The use of data compression, data compaction, and the Use Optimum Block Size (USEOPTBLK) parameter also influence the performance of save operations and restore operations.

As you look at tape drives and their performance rates, you need to understand the tape hardware and the capabilities of that hardware. The different tape drives and IOPs have different capabilities to manipulate data for the best results in their target market. With this in mind, it should be noted that a slower rated tape drive can actually be faster for some customer environments. The following table shows the tape drives that were used in the workload tests and the rates for each drive. (The rates are used later in this document to determine overall performance.)

A study of data compaction was performed on a general sampling of customer data. The study found that compression occurred at a ratio of approximately 2.8 to 1. The performance data for the 2 GB work loads is based on this ratio for tape drives that make full use of the LZ1 compaction algorithm, and the same data compacts at about 1.8 to 1 on drives that do not.

Table 15.0.1		
Drive	Rate (MB/S)	COMPACT
6380	0.3	0.0
6381 #3	0.3	1.8
6382 #3	0.38	1.8
6383	1.5	1.8
6385	1.5	1.8
6386	2.0	1.8
6390	0.5	1.8
7208-342	3.0	1.8
3570	2.2	2.5
3570-С	5.5	2.5
3590 #1	9.0	1.61
3590 Ultra SCSI B model #2	See #2 below	

1. The 3590 does not make use of full compaction. The figures listed here are an attempt to simulate the way it actually works today.

2. The 3590 Ultra SCSI B model tape drive is limited by the AS/400 IOP and Busses. The 3590 Ultra SCSI B model can make use of compaction but is held back by the abilities of the IOP. As an estimate for the 3590 Ultra SCSI B model is to add 20% to the 3590 formulas.

3. The 6381 and 6382 tape drives do not support the system COMPACT parameter. However, data compaction is implemented at the drive hardware level, and is enabled when the *QIC2DC (6381) or the *QIC4DC (6382) density parameters are selected. Use the right tape cartridge and *DEVTYPE for the density parameter on the INZTAP command and these devices will default to the proper QIC density.

15.1 Use Optimum Block Size (USEOPTBLK) parameter

The USEOPTBLK parameter is used to send a larger block of data to tape drives that can take advantage of the larger block size. Every block of data that is sent has a certain amount of overhead that goes with it. This overhead includes block transfer time, IOP overhead, and drive overhead. The block size does not change the IOP overhead and drive overhead, but the number of blocks does. For example, sending 8 small blocks will result in 8 times as much IOP overhead and Drive overhead. With the larger block size, the IOP overhead become less significant. This allows the actual transfer time of the data to become the gating factor. In this example, 8 software operations with 8 hardware operations essentially become 1 software operation with 1 hardware operation when USEOPTBLK(*YES) is specified. The usual result is significantly lower CPU utilization. This also allows the tape device to perform more efficiently.

15.2 Data Compression (DTACPR)

Data compression is the ability to compress strings of identical characters and mark the beginning of the compressed string with a control byte. Strings of blanks from 2 to 63 bytes are compressed to a single byte. Strings of identical characters between 3 and 63 bytes are compressed to 2 bytes. If a string cannot be compressed a control character is still added which will actually expand the data. This parameter is usually used to conserve storage media. If the IOP does not support data compression, the software performs the compression. This situation can require a considerable amount of processing power.

15.3 Data Compaction (COMPACT)

Data compaction is only available at the hardware level. If you wish to use data compaction, the tape drive you choose will need to support it.

15.4 Work Loads

The following work loads were designed to help evaluate the performance of save operations and restore operations. Familiarization with the makeup of the work loads will help clarify the differences in the save rates and restore rates that are presented below.

- NUMX The User Environment workload (NUMX) consists of 4 libraries. The first library contains 4 source files (for a total of 1204 members) that comprise about 39 MB of space. The second library consists of 28 database files, ranging in size from 2 MB to 200 MB, which total 470 MB in size. The third library consists of 200 program objects, with an average size of about 100 KB, for a total size of 20 MB. The fourth library is 12 MB in size and consists of 2156 objects of various types. Therefore, the NUMX workload consists of about 556 MB.
- **NSRC** The NSRC workload consists of the 4 source files that are in the first library of the NUMX. These source files occupy about 39 MB of space and contain a total of 1204 members.
- **2GB** The 2GB workload is a single member database file that is about 2 GB in size. This workload is saved using the SAVOBJ command and restored using the RSTOBJ command.
- **DLO** The DLO workload consists of 8 folders with 3700 documents in the folders. The documents range in size from 53K to 233K with a combined size of 396MB. All of the documents reside in the first level folder structure.
- Integrated File System In the past, the type of data stored in the file system mainly consisted of client programs. Programs don't compact or compress so they are saved or restored at the native rate of the tape drive being used. With the introduction of Lotus Notes and Web functions, more files that contain data are being stored in the file system. With these changes, the rate at which the RST and SAV commands complete has changed because these objects can take advantage of the compaction capabilities of the tape drives.

The following describes save and restore rates that a customer might see depending upon their data and its compaction capabilities. Take a system with an even mixture of client programs, such as, Lotus Notes databases and Web home pages. This example should save and restore in the range of the NUMX workload described in our charts. If the data stored on the system is largely made up of database files, the save/restore rates will probably tend toward the 2GB file type of workload, depending on the size and number of database files. If the data is largely made up of Web files, which tend to be numerous small HTML files such as small home pages, the save rates will tend downward from NUMX toward the NSRC workload.

Web objects can be large images and client databases, just as Lotus Notes database files can be numerous empty or near empty mail files. This would reverse the description above. In all situations the actual data will dictate the save/restore rates and the customer will need to know the type of data they have on their systems in order to estimate the save or restore rates.

15.5 Comparing Performance Data

When comparing the performance data in this document with the actual performance on your system, remember that the performance of save operations and restore operations is data dependent. If the same tape drive was used on data from three different systems, three different rates may result. The performance fluctuation is dependent on the data itself.

The performance of save operations and restore operations is also dependent on the system configuration and the number of DASD units on which the data is stored .

Generally speaking, the large file data that was used in testing for this document was designed to compact at an approximate 2.6:1 ratio. If we were to write a formula to illustrate how performance ratings are obtained, it would be as follows:

((TapeSpeed * LossFromWorkLoadType) * Compaction) = MB/Sec

We would then multiply MB/Sec * 3600 = MB/HR.

But the reality of this formula is that the LossFromWorkLoadType is far more complex than described here. The different work loads have different overheads, different compaction rates, and the drives use different buffer sizes and different compaction algorithms. The attempt here is to group these work loads as examples of what might happen with a certain type of drive and a certain workload.

Note: Remember that these formulas and charts are to give you an idea of what you might achieve from a particular tape drive. Your data is as unique as your company and the correct tape solution must take into account many different factors. These factors include system size, tape drive model, the number of tapes that are required, and whether you are performing an attended or unattended save operation.

Most of the Save/Restore rates listed in this document were obtained from a restricted state measurement. A restricted state measurement is performed when all subsystems are ended using the command ENDSBS SBS(*ALL) so that only the console is allowed to be signed on and running jobs. The new work loads for concurrent and parallel save and restore operations were done on a dedicated system. A dedicated system is one where the system is up and fully functioning but no other users or jobs are running except the save and restore operations. Other subsystems such as QBATCH are required in order to run concurrent and parallel operations. All work loads were deleted before restoring them again.

15.6 Lower Performing Tape Drives

With the lower performing tape drives (6380, 6390), the drives themselves become the gating factor, so the save rates are approximately the same, regardless of system size. For testing purposes on the 6380, compression was substituted for compaction in the formula.

Table 15.6.1 La	Table 15.6.1 Lower performing tape drives LossFromWorkLoadType Approximations (Save Operations)										
Workload Type	e			Amount of Loss							
2GB				95%							
NUMX, DLO				85%							
NSRC			50%								
Example for a	ı 6390	Tape Drive:									
TapeSpeed	*	LossFromWorkLoad	*	Compaction							
.5	*	.95 = (.475)	*	1.8 = (.855) MB/S * 3600 = 3078 MB/HR							

15.7 Medium Performing Tape Drives

The overhead is different for the medium performing tape drives (6385, 3570-B, 6383, 6386). These tape drives have been designed with different technologies and with different markets in mind. It is the differences in these drives that make them difficult to compare. The 3570-B uses optimum block size and LZ1 compaction. Use of USEOPTBLK(*YES) can make the 3570 an efficient drive for systems that are CPU-constrained.

Table 15.6.2 M	Table 15.6.2 Medium performing tape drives LossFromWorkLoadType Approximations (Save Operations)											
Workload Typ	be			Amount of Loss								
2GB				85%								
NUMX, DLO				65%								
NSRC				25%								
Example for	3570 T	ape Drive:										
TapeSpeed	*	LossFromWorkLoad	*	Compaction								
2.2	*	.85 = (1.87)	*	2.5 = (4.68) MB/S * 3600 = 16848 MB/HR								

15.8 Highest Performing Tape Drives

The overhead for the highest performing tape drives (3570-C, 3590, 3590 Ultra SCSI B model) is also different from the other types of tape drives. The high speed tape drives take advantage of optimum block, and high speed data transfer rates. The high speed tape drives are designed to perform best on large files. The use of multiple high speed tape drives concurrently or in parallel can also help to minimize the save window. See section on Multiple Tape Drives for more information.

Table 15.8.1 Higher performing tape drives LossFromWorkLoadType Approximations (Save Operations)							
Workload Type	Amount of Loss						
2GB	95%						
NUMX, DLO	44%						
NSRC	12%						

The 3590 does not make full use of compaction, and the 1.6 is an attempt to simulate the way it actually works today. The 3590 Ultra SCSI B model can make use of compaction but is held back by the abilities of the IOP. As an estimate for the 3590 Ultra SCSI B model add 20% to the 3590 formulas. **Example for 3590 Tape Drive:**

TapeSpeed	*	LossFromWorkLoad	*	Compaction
9.0	*	.95 = (8.6)	*	1.6 = (13.68) MB/S *3600 = 49248 MB/HR

15.9 Multiple Tape Drives

Concurrent Backup and Recovery The ability to save/restore objects from a single library to multiple tape drives or different libraries to multiple tape drives at the same time from different jobs. The work loads that were used for the V4R4 testing were Large file and NUMX. For the large file test, two libraries were created with a 4 GB file in each library. For the NUMX workload, all objects from the 4 NUMX libraries were combined into one library and all of the objects were then duplicated until the library was 6 GB. Then the library was duplicated for the concurrent test. This was different from the V4R3 testing where the objects were duplicated into a single library so that contention for the library would take place.

Parallel Backup and Recovery The ability to save/restore a single object or library across multiple tape drives from the same job. Understand that the function was designed to help those customers, with very large files which are dominating the backup and recovery window. The goal is to provide them with options to help reduce that window. Large objects using multiple tape drives using the parallel function, can greatly reduce the time needed for the object operation to complete as compared to a serial operation on the same object.

Concurrent operations to multiple drives will probably be a better solution for most customers. The customers will have to weigh the benefits of using parallel verses concurrent operations for multiple tape drives in their environment. The following are some thoughts on possible solutions to backup and recovery situations:

- For backup and recovery with a User Mix and small to medium file workloads, the use of concurrent operations will allow multiple objects to be processed at the same time from different jobs, making better use of the tape drives and the AS/400.

- For systems with a lot of data and a few very large files, a mixture of concurrent and parallel might be helpful. (Example: Save all of the libraries to one tape drive, omitting the large files. At the same time run a parallel save of those large files to two or more additional tape drives.)

- For a system dominated by one large file the only way to make use of multiple tape drives is by using the parallel function.

- For systems with a few very large files that can be balanced over the tape drives, use concurrent saves. - Backups where your libraries increase or decrease in size significantly throwing your concurrent saves out of balance constantly, the customer might benefit from the parallel function as the libraries would tend to be balanced against the tape drives no matter how the libraries change. Again this depends upon the size and number of data objects on your system.

- Customers planning for future database growth where they would be adding tape drives in time, might benefit by being able to set up BRMS using *AVAIL for tape drives. Then when a new tape drive is added to the system and recognized by BRMS it will be used, leaving your BRMS configuration the same but benefiting from the additional tape drive. Also the same in reverse, if you lose a tape drive your weekly backup doesn't have to be postponed and your BRMS configuration doesn't need to change, the backup will just use the available tape drives at the time of the save.

The following information and charts are an attempt to help show what might happen in a customer environment with multiple tape drives attached. There are some configuration considerations when using multiple tape devices. A single 2682 bus will allow an Ultra 3590 device to save and restore data at about 60000 MB/HR. If you attempt to add a second drive on that same 2682 bus, you will only get about 1/3 more data using the two tape drives as the 2682 bus will become saturated.

A 2688 bus card is used to attach two towers and if you have a tape drive in each of the towers you will get less data threw than if you have two 2688 bus cards and one tower with a tape drive attached to each 2688 bus card.

A tower sharing dasd and a tape IOP across its bus is not the optimal environment but when looking at multiple tape drives it becomes the only way for large customers to attach all of the dasd and tape drives they need. It is more efficient for a save operation to share the bus with dasd and the tape drive than to put two tape drives on one bus. So rather than having a dasd tower and an IOP tower with two tape drives, use two IOP towers, with a tape IOP and dasd in each of the towers.

For the IBM lab testing, a 12-way system was used with six 2688 bus cards attaching 12 towers. All towers allowed IOP's to be placed in them and each of the towers contained 16 - 8 GB dasd arms. A mixture of 6501 and 6534 tape IOP's were used. Tape drives were used in sequential order for the different tests(i.e. One tape drive test used TAP01, two tape drive test used TAP01 and TAP02, three tape drive test used TAP01, TAP02 and TAP03, etc.). The following shows the tower placement of the tape drives and which towers shared a 2688 bus card:

	2688 2688		2688		2688		2688		2688			
T	ower 1	Tower 2	Tower 3	Tower 4	Tower 5	Tower 6	Tower 7	Tower 8	Tower 9	Tower 10	Tower 11	Tower 12
Т	CAP02	TAP08	TAP04	TAP10	TAP06	TAP07	TAP05		TAP09	TAP03		TAP01

The main bus on the 12-way AS/400 650-2189 will saturate at about 345000 MB/HR. So once you have reached that point adding more tape drives will not change the MB/HR it will just distribute it to more tape drives. As you can also see from the testing the user mix environment (NUMX 6 GB) saturates the system bus at a lower level than the large file, but the maximum number of useful tape drives was the same for both environments.

V TINT Lates	the facts for concurrent save and restore operations are from separate instances												
Table 15.9.1	Table 15.9.1 Concurrent Rates from an AS/400 650-2189 using the 3590 Ultra SCSI B model												
4 GB File Wa	4 GB File Workload using SAVLIB												
# Tape	1	2	3	4	5	6	7	8	9	10			
Drives													
SAVE	60000	118000	175000	225000	280000	320000	325000	345000	345000	345000			
RESTORE	59000	110000	150000	165000	170000	185000	185000	185000	185000	185000			

V4R4 rates for Concurrent save and restore operations are from separate libraries

Table 15.9.2 Concurrent Rates from an AS/400 650 -2189 using the 3590 Ultra SCSI B model NUMX 6 GB Workload using SAVLIB										
# Tape	e 1 2 3 4 5 6 7 8 9 10									
Drives										
SAVE	53000	100000	140000	185000	205000	245000	245000	260000	260000	260000
RESTORE	30000	52000	61000	70000	75000	75000	80000	85000		

V4R4 rates for parallel save and restore operations:

Table 15.9.3 Parallel Rates from an AS/400 Model 650-2189 using the 3590 Ultra SCSI B model 100 GB File Workload using SAVLIBBRM											
# Tape	1 2 3 4 5 6 7 8 9 10										
Drives											
SAVE	60000	93000	135000	195000	240000	269000	297000	315000	345000	345000	
RESTORE	59000	92000	121000	165000	185000	185000	185000	185000	185000	185000	

Other Parallel Runs for comparison

		-											
Table 15.9.4 Paralle	Table 15.9.4 Parallel Rates from an AS/400 Model 740-2070 feature 1513 using the 3590 Ultra SCSI B model												
	1 Tape Drive		2 Tape Drive		3 Tape	e Drive	4 Tape Drive						
	Save	Restore	Save	Restore	Save	Restore	Save	Restore					
4 GB FILE	55000	55000	65000	66000	80000	82000	80000	82000					
16 GB FILE	60000	59000	92000	90000	122000	120000	145000	140000					
8 - 2 GB FILES	60000	59000	93000	92000	126000	120000	145000	145000					
36 GB FILE	60000	59000	93000	90000	133000	126000	160000	160000					
100 GB FILE	60000	59000	93000	92000	135000	121000	182000	168000					
NUMX 6 GB	51000	35000	65000	46000	76000	52000	82000	55000					
NUMX 15 GB	54000	35000	73000	46000	90000	52000	111000	55000					

NOTE: These are the parallel rates that resulted from tests done in the Rochester Lab. In the examples the data used compacted at about 2.5:1 In all of the examples, BRMS was used with either SAVOBJBRM or SAVLIBBRM. All of these examples are meant to help you understand what might happen, not necessarily what will happen, in your environment.

15.10 V4R3 Save and Restore Rates

The save/restore rates that are quoted below are expressed in terms of megabytes per hour (MB/HR). All of these measurements used the default settings in the save and restore commands for compression DTACPR (*DEV), compaction COMPACT(*DEV), and use optimum block USEOPTBLK(*YES).

Note: PLEASE READ THE OTHER SECTIONS TO UNDERSTAND HOW THESE NUMBERS WERE ACHIEVED.

Table 15.10.1 Rate	es from an AS	/400 Model 5	10 Feature 2	144				
				Wor	kload			
	NS	RC	NU	MX	20	βB	DI	LO
Tape Device	Save	Restore	Save	Restore	Save	Restore	Save	Restore
6380	1410	1120	1760	1680	1900	1894	1780	1730
6382	1830	1350	3560	2010	4050	2180	3640	2690
6390	1280	1560	2470	2400	2625	2700	3310	3170
7208-342	2240	1980	9970	8220	15580	15850	14790	4540
6385	1320	1470	5300	2390	9640	9530	1980	2260
3570-В	4500	2400	11800	11200	20600	14300	12400	4650
3570-С	4170	2800	18600	9000	35900	20500	20400	4700
3590	6100	3000	18600	14000	44750	40500	24300	5900

Table 15.10.2 Rates from an AS/400 Model S20 Feature 2166								
	Workload							
	NSRC NUMX 2GB DLO							
Tape Device	Save	Restore	Save	Restore	Save	Restore	Save	Restore
3570-В	5800	4500	14100	14400	20800	21000	15000	10500
3570-С	4648	6000	21900	22800	41500	41000	20200	17800
3590	6200	4690	24000	23900	49800	49900	25000	20000

15.11 V4R4 Rates On New Devices

Table 15.11.1 Rates from an AS/400 Model 170 Feature 2291								
	Workload							
	NSRC NUMX			2GB		DLO		
Tape Device	Save	Restore	Save	Restore	Save	Restore	Save	Restore
6383	2500	2500	6500	6500	9000	9000	9000	3000
6386	2700	2700	8000	8000	12000	12000	11000	3200

Table 15.11.2 Rates from an AS/400 Model 740-2070 feature 1513										
	Workload									
	NSRC NUMX 2GB 16GB DLO						LO			
Tape Device	Save	Restore	Save	Restore	Save	Restore	Save	Restore	Save	Restore
6383	3000	3000	7500	7500	9500	9500			9000	6000
6386	3000	3000	9500	9500	12000	12000			11000	6000
3590 Ultra SCSI	8000	7000	35000	30000	58000	58000	61000	60000	41000	14000
B model #1										

1. The 3590 Ultra SCSI B model tape drive needs the 6534 or the 6501 to achieve the numbers we have published here.

The 2729 IOP doesn't allow the throughput to drive this tape drive any faster than the previous model of 3590. For the 2729 IOP, refer to the V4R3 3590 tape rates to help estimate your expected speeds.

15.12 Save/Restore Rates for Optical Device

The following save and restore performance measurements were made on the 3995 Model C48 Optical Library with a 4X drive. The save and restore rates are expressed in terms of megabytes per hour (MB/HR). They include the time needed to complete the save or restore operation, but not the time that is required for the autochanger to load or unload the optical cartridge. All of these measurements used the default settings in the save and restore commands for compression DTACPR(*DEV), compaction COMPACT(*DEV), and use optimum block USEOPTBLK(*YES).

Table 15.12.1						
Operation	Worl	kload				
	NUMX	2GB				
Save	1925	2160				
Restore	3090	6470				

The following save and restore performance measurements were made on the 3995 model C46 Optical Library with an 8X drive.

Operation Workload				
	NUMX	2GB		
Save	3200	3500		
Restore	3800	7600		

15.13 Hierarchical Storage Management

Hierarchical Storage Management (HSM) provides an automatic way of managing and distributing data between the different storage layers to meet the users needs for accessing data while minimizing the overall cost. The concept of HSM involves the placement of data items in such a way as to minimize the cost of storing your data, while maximizing the accessibility.

The following examples were run on a 20S - 2166 System containing 16 - 6607 DASD units in each of the ASPs. One of the ASPs was compressed and one was not. The rate at which the data was demoted and promoted was approximately the same making HSM the gating factor for performance in this example. The number and type of DASD can and will effect customer performance. See the DASD section in this guide for information on DASD types and the effects of compression.

Table 15.13.1 HSM Rates for Migrating Data Between ASPs	
Workload	Transfer Rate (MB/Hr)
NUMX	5700
2GB	10000
DLO	2200

The rates are expressed in megabytes per hour (MB/HR).
15.14 Save/Restore Tips for Better Performance

- 1. As noted in table 15.0.1 the 6381 and 6382 tape drives are effected by the tape cartridge type and density. For most tape drives the right tape cartridge and density can greatly effect the capacity and speed of your save operation. Some devices look at the cartridge density to decide if the device can use compaction. **USE THE RIGHT TAPE CARTRIDGE FOR YOUR TAPE DRIVE.**
- 2. Using the default setting for the USEOPTBLK parameter of *YES on save commands can significantly improve performance on newer tape drives. This is especially true where the system's CPU is subjected to a heavy workload. A description of this parameter is in one of the first sections of this chapter.
- 3. The placement of the tape IOPs can be important to save and restore performance. The tape IOP should be placed on the bus that has the least amount of DASD attached to it. This is to avoid collisions with the data that is being gathered and the data that is being sent to the tape drive. When using multiple tape drives be sure that your tape IOPs are on separate busses. The 6501 has two ports but only one can be used for parallel or concurrent save or restore operations. So if you have 2 tape drives you will need two 6501 tape IOP cards placed in separate IOP towers.
- 4. In the IBM Rochester lab testing we found that with the speed of the newer 8 GB DASD arms we only need 10 to 16 8 GB DASD as compared to 32 4 GB DASD arms to have the data spread well enough to feed a single 3590 tape drive.
- 5. In the IBM Rochester lab testing we found that we only needed 1 CPU for each tape drive used when running the large file workload. But we needed 1& ¹/₂ CPU's to be able to feed the NUMX data through to each tape drive.
- 6. The optical cables that connect the busses come in different lengths. The 6 to 20 meter cables do not display much difference in affecting the performance of save and restore operations. In the IBM Rochester lab testing determined that save operations ran 55% slower over a 2 kilometer cable as compared to a 6 meter cable. The 2 kilometer cable is only supported on a 266 bus. The 266 bus normally drives the 3590 tape device at around 80% of its capability. When the 2 kilometer cable is added, 3590 tape device is being driven at around 40% of its capability. The 500 meter cable is supported on the 1063 bus. The difference in performance for a 500 meter cable across the 1063 bus using a 3590 tape drive is approximately 10% to 15% less than using a 6 meter cable.
- 7. A tape management system such as BRMS/400 is recommended to keep track of the data and make the most of multiple tape drives.

15.15 New For V4R4

- **Parallel Save and Restore operations:** This function gives the customer the ability to split a single object or library across multiple tape drives using the same job. See the section on Multiple Tape Drives for more information on this topic.
- **3590 Ultra SCSI B model Tape Drive:** In general we saw about a 20% improvement over the previous 3590 for large file save and restore operations on the AS/400. This tape drive needs the 6534 or the 6501 to achieve the numbers we have published in the section on V4R4 rates. The 2729 IOP doesn't allow the throughput to drive this tape drive any faster than the previous model of 3590. For the 2729 IOP refer to the V4R3 3590 tape rates to help estimate your save and restore operations.
- 6383 & 6386 Tape drives: New ¹/₄ inch tape device.
- **3570 & 3590 small object PTF:** V4R4 offers potential performance gains for users with 3570 and 3590 tape devices when backing up a large number of small objects, typically several gigabytes of data with files containing thousands of members. In the IBM Rochester lab performance benchmarks yielded a 10% improvement with a NUMX type workload adjusted to reflect a system with these types of objects. At the extreme end with optimal data we have measured a 66% improvement. We are working with users to get feedback on realistic benefits from this change. In the meantime this change can benefit any user running V4R1 and later by applying the appropriate PTF listed below.

PTFs for previous releases:

- ✤ V4R3 SF54071
- ✤ V4R2 SF54070
- ♦ V4R1.4 SF54069
 ♦ V4R1 SF54068

Chapter 16 IPL Performance

Performance information for Initial Program Load(IPL) is included in this section.

The primary focus of this section is to present data that compares V4R3 IPL times versus V4R4 IPL times across two hardware configurations. The data for both normal and abnormal IPLs is broken down into phases, making it easier to see the detail.

NOTE: The information that follows is based on performance measurements and analysis done in the AS/400 Division laboratory. Actual performance may vary significantly from what is provided here.

<u>16.1 IPL Performance Considerations</u>

There are many factors that affect IPL performance. The wide variety of hardware configurations and software environments available to an AS/400 customer make it difficult to characterize a 'typical' IPL environment and predict the results. The following is a simple description of the IPL tests performed and documented here.

16.2 IPL Benchmark Description

Normal IPL

- Power On IPL (cold start)
- For normal IPLs, benchmark time is measured from power-on to console sign-on screen

Abnormal IPL

- System abnormally terminated causing recovery processing to be done during the IPL. The amount of processing is determined by the activity and reason the system is IPLing.
- For abnormal IPLs, the benchmark consists of bringing up a database workload and letting it run until the desired number of jobs are running on the system. Once the workload is stabilized, a function 22 is issued to the op panel, forcing a main store dump (MSD). This environment simulates a loop or hang situation. The dump is then copied to DASD via the Auto Copy function. The Auto Copy function is enabled through System Service Tools (SST). Once the dump is copied, the system completes the remaining IPL with no user intervention. This is possible using the Auto Copy function, and by putting the key in normal mode shortly after the function 22 is requested. Benchmark time is measured from the time the function 22 is issued, to the time the console sign on screen appears.
- Settings: on the CHGIPLA command the parameter, HDWDIAG, set to (*MIN). All physical files are explicitly journaled. Also logical files are journaled using SMAPP (System Managed Access Path Protection) by using the EDTRCYAP command set to *MIN.

NOTE: Due to some long running tasks (like TCP/IP), all workstations may not be up and ready at the same time as the console workstation displays a sign-on screen.

Large System Benchmark Information

Hardware Configuration

740-2070-1513(12-way) with 40 GB Main Storage DASD / 110 arms, 640 GB of (2, 4, 8 and 16 GB arms) All DASD was RAID protected 2 ASP's defined

Software Configuration

90,000 spool files (30,000 completed jobs with 3 spool files each)
1000 jobs waiting on job queues (active)
9000 delayed jobs waiting on job queues (inactive)
200 remote printers
6000 user profiles
3000 libraries

Database:

- 25 libraries with 2600 physical files and 452 logical files
- 2 libraries with 10,000 physical files and 200 logical files

NOTE:

- Physical files are explicitly journaled
- Logical files are journaled using SMAPP set to *MIN
- Commitment Control used on 20% of the files

Small System Benchmark Information

Hardware Configuration

170-2291 with 128 MB Main Storage DASD / 4 arms, 32 GB (8 GB arms), RAID Protected

Software Configuration

2,000 spool files (2,000 completed jobs with 1 spool file per job) 350 jobs in job queues 100 user profiles 200 libraries

Database:

- 1 library with 100 physical files and 20 logical files
- 1 library with 50 physical files and 10 logical files

<u>16.3 IPL Performance Measurements</u>

The following tables provide a comparison summary of the measured performance data for normal and abnormal IPLs. Results presented do not represent any particular customer environment.

Table 16.3.1 AS/400 Normal IPL Benchmark Summary - Power-On (Cold Start)					
	Large (V4R3 Syste V4R4 System	12-way) m 650-2189 740-2070-1513	Small 3 170-	System 2291	
	V4R3	V4R4	V4R3	V4R4	
Hardware	10:20	10:15	4:50	2:28	
SLIC	1:49	4:35	0:54	1:31	
OS/400	6:02	4:59	2:07	3:28	
Total	18:11	19:49	7:51	7:27	

Measurement units are in minutes and seconds

Generally, the hardware phase is composed of C1xx xxxx and C3xx xxxx SRCs, SLIC is composed of C600 xxxx SRCs, and OS/400 is composed of C900 xxxx SRCs plus time to console sign-on

Table 16.3.2 AS/400 Abnormal IPL Benchmark Summary						
	Large (V4R3 Syste	12-way) em 650-2189	Small System 170-2291			
	V4R4 System	740-2070-1513				
	V4R3	V4R4	V4R3	V4R4		
Processor MSD	10:27	12:12	2:01	2:31		
Hardware IPL	1:46	1:31	0:50	:36		
SLIC MSD IPL with	54:50	42:46	3:02	3:27		
Сору.						
Shutdown	:17	:33	:18	:41		
Hardware re-ipl	7:59	1:58	3:12	:48		
SLIC re-ipl	6:38	5:53	1:36	1:30		
OS/400	23:33	23:19	5:26	4:51		
Total	1.45.19	1.28.12	16.26	14.24		

Measurement units are in hours, minutes and seconds.

MSD is Main Store Dump. General IPL phase as it relates to the SRCs posted on the operation panel: Processor MSD includes the C1xx xxxx and D1xx xxxx right after the function 22 is issued. Hardware IPL is the next phase which includes the following group of C1xx xxxx and C3xx xxxx SRCs. SLIC MSD IPL with Copy follows with the next series of C6xx xxxx. The copy occurs during the C6xx 4404 SRCs. Shutdown includes the Dxxx xxxx SRCs. Hardware re-ipl includes the next phase of C1xx xxxx and C3xx xxxx. SLIC re-ipl follows which are the C600 xxxx SRCs. OS/400 completes with the C900 xxxx SRCs plus the time to console sign-on.

ACTUAL CUSTOMER TIMES MAY VARY SIGNIFICANTLY FROM THIS.

16.4 MSD Affects on IPL Performance Measurements

Through some experimental testing we have found that the time spent in MSD copying the data to disk is related to the number of dasd arms available. The following are times with different dasd arms available. These timings are for the C6xx 4404 SRC portion of the MSD, not the entire time spent doing the MSD portion of the IPL. C6xx 4404 is the time during the MSD where main store is copied to the dasd. By understanding your system configuration, this information and the other information in this document, can help you estimate the amount of time your system may take to IPL when a main storage dump is needed or happens.

The system used for this test was a 740 270-1513 with 40 GB main storage and 200 8 GB dasd arms all RAID protected. The number of arms refers to the number of arms in the ASP where MSD was copied.

Table 16.4.1	10 Arms	20 Arms	36 Arms	64 Arms	80 Arms	112 Arms	200 Arms
40 GB MSD Copy (C600-4404)	2 hr 09 min	1 hr 50 min	1 hr 07 min	34 hr	30 min	22 min	13 min

16.5 IPL Tips

Although IPL duration is highly dependent on hardware and software configuration, there are tasks that can be performed to reduce the amount of time required for the system to perform an IPL. The following is a partial list of recommendations for IPL performance:

- Reduce the number of jobs on the system. The best way to do this is by removing unnecessary spool files. Use the Display Job Tables (DSPJOBTBL) command to monitor the size of the job table(s) on the system. The Change IPL Attributes (CHGIPLA) command is used to compress job tables when there are a large amount of available job table entries.
- Reduce the number of device descriptions remove any obsolete device descriptions from the system.
- Control the level of hardware diagnostics that are run during an IPL. On the CHGIPLA command, by specifying the default, HDWDIAG(*MIN), the system will perform only a minimum, critical set of hardware diagnostics. This type of IPL is appropriate in most cases. The exceptions include a suspected hardware problem, or when new hardware, such as additional memory, is being introduced to the system.
- Reduce the amount of rebuild time for access paths during an IPL by using System Managed Access Path Protection (SMAPP). The AS/400 Backup and Recovery book (SC41-5304) describes this method for protecting access paths from long recovery times during an IPL.
- For additional information on how to improve IPL performance, refer to AS/400 Basic System Operation, Administration, and Problem Handling (SC41-5206) - or to the redbook The System Administrator's Companion to AS/400 Availability and Recovery (SG24-2161).

Chapter 17. Integrated Netfinity Server

This chapter presents performance results for the Integrated Netfinity Server on the AS/400. In addition to results, this paper contains tips for maximizing Integrated Netfinity Server performance and performance monitoring techniques for both the AS/400 and Integrated Netfinity Server. Also included in this paper is an introduction to the Integrated Netfinity Server and an overview of NetBench, a popular PC server industry benchmark.

17.1 Introduction

The Integrated Netfinity Server extends the utility of the AS/400 by combining a PC server with Windows NT inside the AS/400. There are two version of the Integrated Netfinity Server: one PCI based version for AS/400e series models; and a SPD 'book package' version for AS/400 Advanced Series models or SPD integrated expansion units attached to AS/400e series models, as shown in Figure 17.1.

The Integrated Netfinity Server has an 333Mhz Intel Pentium II processor with 512 KB L2 Cache (L2 Cache runs at core processor speed), serial and parallel ports, and options for 64 Mbytes to 1GBytes of ECC/EDO memory as shown in the table of Figure 17.2.. Three different network cards (16 Mb token-ring, 10 and 10/100 Mbps Ethernet) are supported. A monitor, keyboard and mouse must be attached to the card to act as an NT console. Windows NT device drivers are provided to share the AS/400's disk, tape, and CD-ROM drives. Integrated Netfinity Server operations and systems administration are integrated with the AS/400. It also provides a platform for integrated applications between the AS/400 and Windows NT.



Figure 17.1. PCI and SPD Bus Versions of the Integrated Netfinity Server

	Integrated PC	Integrated PC		
	Server	Server		
	PCI	SPD		
Processor	Pentium II	Pentium II		
	333 Mhz	333 Mhz		
L2 Cache	512KB	512KB		
	(at proc speed)	(at proc speed)		
Memory	Up to 1 GB	Up to 1 GB		
AS/400	AS/400e series	AS/400e series or		
	with	Advanced Series		
	PCI Bus	with SPD Bus		
AS/400 Slots	Pre-reserved	3 SPD		
LAN	1-2	1-3		
Adapters	Token-Ring	Token-Ring		
	Ethernet 10/100	Ethernet 10/100		
Device	Parallel Port	Parallel Port		
Options	1 Serial Port	2 Serial Ports		
Software	NT Server 4.0	NT Server 4.0		
Support	AS/400 Firewall	AS/400 Firewall		



The AS/400 Integrated Netfinity Server runs Microsoft Windows NT Server Version 4.0: the standard CD-ROM version that can be purchased from any Microsoft reseller. The Integrated Netfinity Server has passed the Microsoft Compatibility Tests for Windows NT Server V4.0, signified by the logo display: Designed for Windows NT. For more information see the Microsoft Hardware Compatibility List (HCL) at www.microsoft.com/hwtest/hcl (Category: Misc., Manufacturer: IBM). Microsoft NT Server itself has not been modified to run on the Integrated Netfinity Server. We have provided device drivers to attach the Integrated Netfinity Server and use the AS/400's disk, tape and CD-ROM drives.

17.2 NT Server Benchmark: NetBench 5.01

The NetBench performance benchmark from Ziff-Davis Benchmark Operation (ZDBOp) was used to measure Integrated Netfinity Server performance. The Ziff-Davis Benchmark Operation (ZDBOp) is a large independent developer of benchmark software. ZDBOp is a division of Ziff-Davis Inc.



Figure 17.3. NetBench Results

Netbench was chosen because of its popularity, widespread use, and it is from an independent developer. You can find out more about NetBench on the Internet at: http://www.zdnet.com/zdbop.

NetBench 5.01 is used to measure file serving performance. NetBench measures file serving performance by determining how well a server handles file I/O requests from clients. NetBench mirrors the way leading PC applications perform network file operations on a file server. The applications profiled by NetBench for 32-bit Windows clients are:

- Adobe PageMaker 6.0
- Borland Paradox 7.0
- CorelDRAW! 6.0
- Lotus WordPro 96
- Microsoft Access 7.0
- Microsoft Excel 7.0
- Microsoft PowerPoint 7.0
- Microsoft Word 7.0
- Windows 95 Explorer

Performance was measured using NetBench version 5.01 standard disk mix test suite NBDM_60.TST, shows the results from running the Netbench disk mix on the 333 MHz Integrated Netfinity Server PCI, the 200 MHz Integrated Netfinity Server PCI, and against a Compaq Prolient 1600 with a single 450 MHz

Pentium II. The same number of similar characteristic hard drives were used for the Compaq Server and the AS/400. The Compaq server disks were striped via Windows NT.

The AS/400 used is a Model 170 with a system expansion unit. The PCI bus versions of the Integrated Netfinity Server were used. One test had the Integrated Netfinity Server controlled by the multifunction I/O processor (MFIOP), which also controls the disk drives. The other test had the Integrated Netfinity Server under a separate communications IOP in the side car.

From this chart we can glean a lot of important data. Remember, this benchmark represents the type of file serving demands popular desktop PC applications would place on an NT server. At its optimum point, the 333 MHz Integrated Netfinity Server was able to serve 7.5 million bytes/sec. Notice that after this point (the knee of the curve), the throughput falls off. This happens because the contention for server resources and the overhead of managing clients has finally reached a point where it makes the server throughput decrease. Notice that the decline is gradual. This indicates that the Integrated Netfinity Server will perform well (predictably) under a heavy load.

The faster Compaq server achieved 13% faster throughput over the 333Mhz Integrated Netfinity Server, which achieved 39% improvement over the 200Mhz Integrated Netfinity Server. Running the Integrated Netfinity Server under an separate communications IOP verses under the main MFIOP had no appreciable effect on throughput.

NetBench is a synthetic benchmark. Each PC client used in the benchmark represents the load of many "real life" clients. This is done to make the test set up more practical. The load one of these PC's puts on the server is dependent on how fast it is. A Pentium client will put a larger load on the server being tested than a 486 client. For that reason we specify the configuration of the client PC's in 17.6 NetBench Benchmark Details.238 Thus, the number of clients on the x-axis of is <u>NOT</u> an indication of how many "real life" PC clients you could expect to handle with the Integrated Netfinity Server. The peak throughput (indicating the maximum file server bandwidth that the Integrated Netfinity Server can be expected to handle) is the useful metric for determining file server capacity.

17.3 Effects of NetBench on the AS/400

The Integrated Netfinity Server uses AS/400 DASD for its hard drives. This is accomplished by special NT DASD device drivers written for the Integrated Netfinity Server. The NT DASD device drivers cooperate with the AS/400 to perform the DASD operations, so AS/400 CPU resource is used as well. The effect on the AS/400 from running applications on the Integrated Netfinity Server is primarily on the DASD subsystem, secondarily on the AS/400 CPU. Using the AS/400 performance monitor, the load on the AS/400 CPU and DASD subsystem was measured. Figure 17.4 and Figure 17.6 show the CPU loads during the Netbench diskmix test.



Figure 17.4 AS/400 Performance with Server under Comm IOP

Operating System	V4R3
Model	170-2292
CPU Performance	220 CPW
Memory	512 MB
DASD 4	8 GB 6713
MFIOP	675A-003
Comm IOP	2809-001
Integrated Netfinity Server	2850-012

Figure 17.5. AS/400 Test System Configuration

The results in Figure 17.4 show the amount of AS/400 resource used when running NetBench while the Integrated Netfinity Server resides under a separate Comm IOP. The AS/400 CPU utilization ranges from 0.4 - 11.7%. This information is useful when sizing your AS/400 with an Integrated Netfinity Server. The CPU performance for the system used in these benchmark measurements (220 CPW) is given as an aid for comparison to other AS/400 models. CPW stands for Commercial Processing Workload. CPW numbers are given for all AS/400 models in the AS/400 System Handbook (GA19-5486).



Figure 17.6. AS/400 Performance with Server under MFIOP

The results in Figure 17.6 show the AS/400 resource used when running Netbench and the Integrated Netfinity Server resides under the MFIOP. As you can see, the work load has been transferred from the Comm IOP in Figure 17.4 to the MFIOP. The sharing of the MFIOP has not disproportionately increased the MFIOP utilization.

The amount of DASD performance to reserve for the Integrated Netfinity Server requires careful consideration. If your file serving demands are constantly going to be equivalent to the peak NetBench throughput, then reserving 35% of the AS/400's DASD throughput capacity for Integrated Netfinity Server may be necessary. However, a large proportion of server installations will probably be at the average or even somewhat below average, yielding a DASD throughput capacity of less than 17%. Again, please remember that these are "rules of thumb" to get started. Performance monitoring should be used to verify and refine early estimates.

17.4 Performance Tips and Techniques

Optimizing NT Server

When setting up your Windows NT Server, you can optimize how it uses memory. To get to the server optimization window, do the following:

- 1. From the task bar choose Start -> Settings -> Control panel
- 2. Double click the Network icon

V4R4 Performance Capabilities Reference

8 Copyright IBM Corp. 1999

- 3. Choose the Services tab and choose server as show in Figure 17.7.
- 4. Click on the properties button
- 5. Refer to Figure 17.8 to choose the optimum setting for your server

Computer Browse Victiss of Internet I Microsoft, 101 / IP NetBicS Interlate RPO Contiguration RPO Contiguration Workstation	nformation Se Finnting	arver3.0	
Ad-	Beminie	Emperties	La-nte
- Description.			
natallis support for the part of Micropoliting	ne SME (Serv svorking:	er Message Block)	proto coli, the
		6 3	

Figure 17.7. Windows NT Server Network Window

Optimization	CK
C Minimize Mamory Uses	Cance
C dolance	Here
Maximize Throughputfor Fills Sharing	- Te b
C Maximize Throughput for Network Applications	

Figure 17.8. Windows NT Server Optimization Window

Choose one of the optimizations depending on your server use.

Maximize Throughput for File Sharing - Allocates the maximum memory for file sharing applications. **Maximize Throughput for Network Applications** - Optimizes memory for server applications that do their own memory caching. Many database server applications do this.

For the type of serving we are discussing here, *Maximize Throughput for File Sharing* would be the recommended optimization.

Key items to monitor on Integrated Netfinity Server.

CPU utilization, Network throughput, disk utilization and disk throughput are the initial items to monitor on the Integrated Netfinity Server to detect performance bottlenecks. The Windows NT Performance Monitor can be used to monitor these. The performance monitor (as shown in Figure 9) can be used to monitor the following objects: Logical Disk, Memory, Process, Processor, System, and Server. For an overview of Network throughput, the Total bytes/sec component in the Server object can be monitored. Performance data can be logged over several days to create an overview of your server performance.



Figure 17.9. Windows NT Performance Monitor

Key items to monitor on AS/400

The AS/400 Performance Monitor can be used to monitor AS/400 CPU, DASD utilization and throughput. The performance monitor can be run over a period of days to give an overview of the effect of the Integrated Netfinity Server on AS/400 performance. The system and component reports are good sources of this data. The component report gives AS/400 CPU utilization and peak DASD utilization over time, as show in Figure 17.10.

Component Report Component Interval Activity NetBench run 200 MHz IPCS

Member:	NBIPC	S128	Model/	Serial:	170-22	292/10-2N	/IPRM	Main	Storage:	512	.0 M	Started:	08/	06/98
Library.	KKISD		System	Ivaille.	КСПА	or Kivi		Versio	on Release.	4/3.	.0	Stopped.	08/	00/98
							Disl	k I/O	Hi	gh		Pool		Excp
Itv	Tns/	Rsp/	DDM	CP	U Utiliza	tion	Per S	econd	Utiliza	ation]	Faults/Sec-		per
End	Hour	Tns	I/O	Total	Inter	Batch	Sync	Async	Disk	Unit	Mch	User	ID	Second
16:47	47	.00	0	0.4	.0	.0	0.5	0.1	2	0005	0	0	03	14
16:52	0	.00	0	0.3	.0	.0	0.4	0.2	3	0006	0	0	03	3
16:57	0	.00	0	0.4	.0	.0	0.3	0.1	2	0005	0	0	03	3
17:02	0	.00	0	0.4	.0	.0	0.3	0.1	2	0006	0	0	03	3
17:07	0	.00	0	2.3	.0	.0	0.2	0.1	3	0005	0	0	03	3
17:12	0	.00	0	2.0	.0	.0	0.2	0.1	7	0006	0	0	03	3
17:17	0	.00	0	2.1	.0	.0	0.3	0.1	7	0006	0	0	02	3
17:22	0	.00	0	2.0	.0	.0	0.2	0.1	5	0005	0	0	02	3
17:27	0	.00	0	2.0	.0	.0	0.2	0.1	6	0005	0	0	03	3
17:32	0	.00	0	4.4	.0	.0	0.2	0.1	8	0005	0	0	03	3
17:37	0	.00	0	4.1	.0	.0	0.2	0.1	12	0006	0	0	03	3
17:42	0	.00	0	4.1	.0	.0	0.2	0.1	12	0006	0	0	03	3
17:47	0	.00	0	4.2	.0	.0	0.3	0.1	10	0005	0	0	02	3
17:52	0	.00	0	4.2	.0	.0	0.2	0.1	11	0006	0	0	03	3
17:57	0	.00	0	6.5	.0	.0	0.2	0.1	12	0005	0	0	03	3
18:02	0	.00	0	6.3	.0	.0	0.2	0.1	15	0006	0	0	03	3
18:07	0	.00	0	6.1	.0	.0	0.2	0.1	13	0006	0	0	03	3
18:12	0	.00	0	6.0	.0	.0	0.2	0.1	15	0006	0	0	03	3
18:17	0	.00	0	6.1	.0	.0	0.2	0.1	16	0005	0	0	03	3
18:22	0	.00	0	6.0	.0	.0	0.2	0.1	15	0005	0	0	03	3
18:27	0	.00	0	7.6	.0	.0	0.2	0.1	18	0006	0	0	03	3
18:32	0	.00	0	11.7	.0	.0	0.2	0.1	34	0005	0	0	03	3
18:37	0	.00	0	11.7	.0	.0	0.2	0.1	33	0005	0	0	03	3
18:42	0	.00	0	10.8	.0	.0	0.2	0.1	29	0005	0	0	03	3
18:47	0	.00	0	10.9	.0	.0	0.2	0.1	31	0005	0	0	03	3
18:57	0	.00	0	10.1	.0	.0	0.2	0.1	28	0005	0	0	03	3
Itv End			Inter	val end tin	ne (hour a	nd minute	:)							
Tns / Hou	ır		Num	ber of inte	ractive tra	ansactions	per hour							
Rsp / Tns			Aver	age intera	ctive trans	action resp	ponse time	in second	s					
DDM I/O)		Num	ber of DD	Ms									
Total CPU	U Utilizatio	n	Perce	entage of a	vailable C	CPU time u	used by in	teractive ar	nd batch job	s. This is	s the ave	rage of all	proces	sors
Inter CPU	J Utilization	1	Perce	ntage of a	vailable C	CPU time u	used by in	teractive jo	bs. This is	the avera	ge of all	processors	-	
Batch CP	U Utilizatic	on	Perce	entage of a	vailable (CPU time i	used by ba	tch jobs. 1	This is the a	verage of	all proc	essors		

Percentage of available CPU time used by batch jobs. This is the average of all processors ---

---Average synchronous disk I/O operations per second

Average asynchronous disk I/O operations per second ---

--Percent of utilization of the most utioized disk arm during this interval

Disk arm which had the most utilization during this interval ---

Figure 17.10.	AS/400	Performance	Monitor	Component	Report
---------------	--------	-------------	---------	-----------	--------

17.5 Summary

Sync Disk I/O Per Sec

High Disk Utilization

High Utilization Unit

Async Disk I/O Per Sec

The AS/400 Integrated Netfinity Server with Windows NT Server V 4.0 is a full NT file, print and application server. It provides flexibility for AS/400 applications and NT services in a combination server with improved hardware control, availability, and reduced maintenance costs.

The Integrated Netfinity Server performs well as a file server for popular PC applications, using the AS/400 DASD for its hard drive. As part of the preparation for a combination server installation, care should be taken to estimate the expected workload of the NT server and reserve AS/400 resources for the Integrated Netfinity Server.

17.6 NetBench Benchmark Details

AS400 Server Configuration

AS/400 Machine	9406-170 w sidecar
OpSys	AS400 Version V4R3 - GA Cum Pkg #5
Processor	#2292 (220/30 CPW)
Memory	512MB (#3302x4)
MFIOP	675A-003
Storage Ctl	2740-001 (PCI Raid with soft write cache)
Comm IOP	2809-001
Disks	4 x 6713-050 (8.58G) 7200 rpm, Ultra-Wide SCSI.
BUS	PCI

333 MHz Integrated Netfinity Server

Integrated Netfinity Server	2850-012 (333 Mhz) (under Comm IOP or MFIOP)
OpSys	Windows NT Server 4.0, SP4
Processor	333 Mhz Pentium II
Lvl2 Cache	512KB (at core processor speed)
Memory	512MB ECC/EDO RAM (66 Mhz memory bus)
Disk Driver	qvndvsdd.sys
Disk Drives	Test share is on single virtual drive.
BUS	PCI
Network Adaptor	2 x 2838 10/100 Ethernet PCI Cards
Network Driver	pcntn4m.sys v4.00.005 (AMD PCNet Family)

200 MHz Integrated Netfinity Server

Integrated Netfinity Server	2850-011 (200 Mhz) (under Comm IOP)
OpSys	Windows NT Server 4.0, SP4
Processor	200 Mhz Pentium Pro
Lvl2 Cache	256KB (at core processor speed)
Memory	512MB ECC/EDO RAM (66 Mhz memory bus)
Disk Driver	qvndvsdd.sys
Disk Drives	Test share is on single virtual drive.
BUS	PCI
Network Adaptor	2 x 2838 10/100 Ethernet PCI Cards
Network Driver	pcntn4m.sys v4.00.005 (AMD PCNet Family)

Compaq Server

Machine	Compaq Prolient 1600 - Series 4070				
OpSys	Vindows NT Server 4.0, SP4				
Processor	450 Mhz Pentium II				
Lvl2 Cache	512KB				
Memory	512MB SDRAM				
Drive Adaptor	Ultra-Wide SCSI 3				
Disk Driver	Symc810.sys v4.00				
Disk Drives	4 x Compaq Ultra-Wide SCSI 3, 7200RPM,				
	4Gig, Hot Pluggable, part #272577-001				
NTFS Logical Drive Striped across the 4 drives.					
	(The system partition cannot be striped).				
Network Adaptor	1 x Compaq Dual Port 10/100 TX UTP				
Network Driver	Netflx3.sys v4.24				

Netbench Controller and Clients

	Controller	Clients
Machine	IBM PC 300PL	IBM PC 300PL (59 Units)
OpSys	Windows NT Server 4.0, SP4	Windows NT Workstation 4.0, SP3
Processor	400 Mhz Pentium II	400 Mhz Pentium II
Lvl2 Cache	512KB	512KB
Memory	64MB SDRAM / 100Mhz bus	64MB SDRAM / 100Mhz bus
Disk Driver	Intel piixide.sys 2.03.0, 3/10/98	Intel piixide.sys 2.03.0, 3/10/98
Disk Drives	1 x 6.4GB IDE (SMART Ultra ATA)	1 x 6.4GB IDE (SMART Ultra ATA)
Network Adaptor	IBM 10/100 Etherjet PCI	IBM 10/100 Etherjet PCI
Network Driver	ibmfentsys v 3.00.07	ibmfentsys v 3.00.07

Network Configuration

Network	100Mbps Ethernet - Full Duplex
Clients	59
Switches	8 x 12 Port 3Com 10/100 SuperStack Switches (3C16464A),
Network Segments	2

NetBench 5.01 Test bed:

The NetBench measurements were conducted using Ziff-Davis' NetBench 5.01 running the Disk Mix with Windows NT Workstation 4.0 as described below:

Version: NetBench 5.01

Mixes:

- Disk Mix (NBDM_60.TST, later mix times increased to allow more test iterations).
- Clients 1,4,8,12,16,20,24,28 Mix runtime = 11 Minutes
- Clients 32,36,40,44,48,52,56,59 Mix runtime = 16 Minutes
- (60th client was unavailable)
- Client workspace: Default
- Ramp up and down: 30 seconds

Network Operating System: Microsoft windows NT Server 4.0

V4R4 Performance Capabilities Reference 8 Copyright IBM Corp. 1999 All products used for these measurements are shipping versions available to the general public. All measurements were performed without independent verification by Ziff-Davis.

17.7 Additional Sources of Information

Integrated Netfinity Server URL: <u>http://www.as400.ibm.com/nt</u> Microsoft Hardware Compatibility Test URL: <u>http://www.microsoft.com/hwtest/hcl</u> (Cat:MISC Co:IBM)

Redbook: "AS/400 - Running Windows NT on the Integrated Netfinity Server SG24-2164" *Manual:* "OS/400 - AS/400 Integration with Windows NT Server SC41-5439"

Chapter 18. Logical Partitioning (LPAR)

18.1 Introduction

Logical partitioning (LPAR) is a mode of machine operation where multiple copies of operating systems run on a single physical machine.

A *logical partition* is a collection of machine resources that are capable of running an operating system. The resources include processors (and associated caches), main storage, and I/O devices. Partitions operate independently and are logically isolated from other partitions. Communication between partitions is achieved through I/O operations.

The *primary partition* provides functions on which all other partitions are dependent. Any partition that is not a primary partition is a *secondary partition*. A secondary partition can perform an IPL, can be powered off, can dump main storage, and can have PTFs applied independently of the other partitions on the physical machine. The primary partition may affect the secondary partitions when activities occur that cause the primary partition's operation to end. An example is when the PWRDWNSYS command is run on a primary partition. Without the primary partition's continued operation all secondary partitions are ended.

18.2 Considerations

This section provides some guidelines to be used when sizing partitions versus stand-alone systems. The actual results measured on a partitioned system will vary greatly with the workloads used, relative sizes, and how each partition is utilized. For information about CPW values, refer to *Appendix D*, "*AS/400 CPW Values*".

When comparing the performance of a standalone system against a single logical partition with similar machine resources, do not expect them to have identical performance values as there is LPAR overhead incurred in managing each partition. For example, consider the measurements we ran on a 4-way system using the standard AS/400 Commercial Processing Workload (CPW) as shown in the chart below.

For the standalone 4-way system we used we measured a CPW value of 1950. We then partitioned the standalone 4-way system into two 2-way partitions. When we added up the partitioned 2-way values as shown below we got a total CPW value of 2044. This is a 5% increase from our measured standalone 4-way CPW value of 1950. I.e. (2044-1950)/1950 = 5%. The reason for this increased capacity can be attributed primarily to a reduction in the contention for operating system resources that exist on the standalone 4-way system.

Separately, when you compare the CPW values of a standalone 2-way system to one of the partitions (i.e. one of the two 2-ways), you can get a feel for the LPAR overhead cost. Our test measurement showed a capacity degradation of 3%. That is, two standalone 2-ways have a combined CPW value of 2100. The total CPW values of two 2-ways running on a partitioned four way, as shown above, is 2044. I.e. (2100-2044)/2044 = -3%.

The reasons for the LPAR overhead can be attributed to contention for the shared memory bus on a partitioned system, to the aggregate bandwidth of the standalone systems being greater than the bandwidth of the partitioned system, and to a lower number of system resources configured for a system partition than on a standalone system. For example on a standalone 2-way system the main memory available may be X, and on a partitioned system the amount of main storage available for the 2-way partition is X-2.



Figure 18.1. LPAR Performance Measured Against Standalone Systems

In summary, the measurements on the 4-way system indicate that when a workload can be logically split between two systems, using LPAR to configure two systems will result in system capacities that are greater than when the two applications are run on a single system, and somewhat less than splitting the applications to run on two physically separate systems. The amount of these differences will vary depending on the size of the system and the nature of the application.

18.3 Performance on a 12-way system

As the machine size increases we have seen an increase in both the performance of a partitioned system and in the LPAR overhead on the partitioned system. As shown below you will notice that the capacity increase and LPAR overhead is greater on a 12-way system than what was shown above on a 4-way system.

Also note that part of the performance increase of an larger system may have come about because of a reduction in contention within the CPW workload itself. That is, the measurement of the standalone 12-way system required a larger number of users to drive the system's CPU to 70 percent than what is required on a 4-way system. The larger number of users may have increased the CPW workload's internal contention. With a lower number of users required to drive the system's CPU to 70 percent on a standalone 4-way system., there is less opportunity for the workload's internal contention to be a factor in the measurements.

The overall performance of a large system depends greatly on the workload and how well the workload scales to the large system. The overall performance of a large partitioned system is far more complicated because the workload of each partition must be considered as well as how each workload scales to the size of the partition and the resources allocated to the partition in which it is running. While the partitions in a system do not contend for the same main storage, processor, or I/O resources, they all use the same main storage bus to access their data. The total contention on the bus affects the performance of each partition, but the degree of impact to each partition depends on its size and workload.

In order to develop guidelines for partitioned systems, the standard AS/400 Commercial Processing Workload (CPW) was run in several environments to better understand two things. First, how does the sum of the capacity of each partition in a system compare to the capacity of that system running as a single image? This is to show the cost of consolidating systems. Second, how does the capacity of a partition compare to that of an equivalently sized stand-alone system?

The experiments were run on a 12-way 740 model with sufficient main storage and DASD arms so that CPU utilization was the key resource. The following data points were collected:

- Stand-alone CPW runs of a 4-way, 6-way, 8-way, and 12-way
- Total CPW capacity of a system partitioned into an 8-way and a 4-way partition
- Total CPW capacity of a system partitioned into two 6-way partitions
- Total CPW capacity of a system partitioned into three 4-way partitions

The total CPW capacity of a partitioned system is greater than the CPW capacity of the stand-alone 12-way, but the percentage increase is inversely proportional to the size of the largest partition. The CPW workload does not scale linearly with the number of processors. The larger the number of processors, the closer the contention on the main storage bus approached the contention level of the stand-alone 12-way system.

For the partition combinations listed above, the total capacity of the 12-way system increases as shown in the chart below.



Figure 18.2. 12 way LPAR Throughput Example

To illustrate the impact that varying the workload in the partitions has on an LPAR system, the CPW workload was run at an extremely high utilization in the stand-alone 12-way. This high utilization increased the contention on the main storage bus significantly. This same high utilization CPW benchmark was then run concurrently in the three 4-way partitions. In this environment, the total capacity of the partitioned 12-way exceeded that of the stand-alone 12-way by 18% because the total main storage bus contention of the three 4-way partitions is much less than that of a stand-alone 12-way.

The capacity of a partition of a large system was also compared to the capacity of an equally sized stand-alone system. If all the partitions except the partition running the CPW are idle or at low utilization, the capacity of the partition and an equivalent stand-alone system are nearly identical. However, when all of the partitions of the system were running the CPW, then the total contention for the main storage bus has a measurable effect on each of the partitions.

The impact is greater on the smaller partitions than on the larger partitions because the relative increase of the main storage bus contention is more significant in the smaller partitions. For example, the 4-way partition is degraded by 12% when an 8-way partition is also running the CPW, but the 8-way partition is only degraded by 9%. The two 6-way partitions and three 4-way partitions are all degraded by about 8% when they run CPW together. The impact to each partition is directly proportional to the size of the largest partition.

18.4 LPAR Measurements

The following chart shows measurements taken on a partitioned 12-way system with the system's CPU utilized at 70 percent capacity. The system was at the V4R4M0 release level.

Note that the standalone 12-way CPW value of 4700 in our measurement is higher than the published V4R3M0 CPW value of 4550. This is because there was a contention point that existed in the CPW workload when the workload was run on large systems. This contention point was relieved in V4R4M0 and this allowed the CPW value to be improved and be more representative of a customer workload when the workload is run on large systems.

Table 18.1 12-way system measurements							
ΙΡΔR	Stand alone	Total	CPW		LPAR CPW	T	Average I PAR
Configuration	12-way CPW	LPAR CPW	Increase	Primary	Secondary	Secondary	Overhead
8-way, 4-way	4700	5020	7%	3330	1690	n/a	10 %
(2) 6-ways	4700	5140	9%	2605	2535	n/a	9 %
(3) 4-ways	4700	5290	13%	1770	1770	1750	9 %

While we saw performance improvements on a 12-way system as shown above, part of those improvements may have come about because of a reduction in contention within the CPW workload itself. That is, the measurement of the standalone 12-way system required a larger number of users to drive the system's CPU to 70 percent than what is required on a 4-way system. The larger number of users may have increased the CPW workload's internal contention.

With a lower number of users required to drive the system's CPU to 70 percent on a standalone 4-way system, there is less opportunity for the workload's internal contention to be a factor in the measurements.

The following chart shows our 4-way measurements.

Table 18.2 4-way system measurements						
ΙΡΛΡ	Stand alone	Total	otal PAR PW Increase	LPAR CPW		
Configuration	4-way CPW	LPAR CPW		Primary	Secondary	Average LPAR Overhead
(2) 2-ways	1950	2044	5%	1025	1019	3 %

The following chart shows the overhead on n-ways of running a single LPAR partition alone vs. running with other partitions. The differing values for managing partitions is due to the size of the memory nest and the number of processors to manage (n-way size).

Table 18.3 LPAR overhead per partition					
Processors	Measured	Projected			
2	-	1.5 %			
4	3.0 %	-			
8	-	6.0 %			
12	9.0 %	-			

The following chart shows projected LPAR capacities for several LPAR configurations. The projections are based on measurements on 1 and 2 way measurements when the system's CPU was utilized at 70 percent capacity. The LPAR overhead was also factored into the projections. The system was at the V4R4M0 release level.

Table 18.4 Projected LPAR Capacities						
LPAR Configuration Projected CPW Increase						
Number	Processors	FIOJECIEU LFAK CFW	Over a Standalone 12-way			
12	1-ways	5920	26 %			
6	2-ways	5700	21 %			

18.5 Summary

On a partitioned system the capacity increases will range from 5% to 26%. The capacity increase will depend on the number of processors partitioned and on the number of partitions. In general the greater the number of partitions the greater the capacity increase.

When consolidating systems, a reasonable and safe guideline is that a partition may have about 10% less capacity than an equivalent stand-alone system if all partitions will be running their peak loads concurrently. This cross-partition contention is significant enough that the system operator of a partitioned system should consider staggering peak workloads (such as batch windows) as much as possible.

Chapter 19. Miscellaneous Performance Information

19.1 Public Benchmarks (TPC-C, SAP, RPMark, NotesBench)

TPC-C Commercial Performance

The Transaction Processing Performance Council's TPC Benchmark C (TPC-C (**)) is a public benchmark that stresses systems in a full integrity transaction processing environment. It is more closely related to general business computing than prior benchmarks (such as TPC-A and TPC-B), but the functional emphasis may still vary significantly from an actual customer environment. For additional information on the benchmark and current results, please refer to the TPC's web site at http://www.tpc.org.

In addition to referring to the TPC web site, request the 'AS/400 TPC-C Results' package from MKTTOOLS for a set of questions and answers about the AS/400 TPC-C results. To request this package issue the following command on VM:

TOOLCAT MKTTOOLS GET AS4TPCC PACKAGE

Information for can be found at: http://www.tpc.org/bench.results.html

SAP Performance Information

In September, 1995, SAP AG of Walldorf, Germany announced availability of its R/3 suite of client/server business applications to IBM's AS/400 Advanced Series. The pilot phase of R/3 release 3.0 for AS/400 began in fourth quarter '95, with a controlled availability phase beginning in 3/96 and general availability in 7/96. SAP's R/3 suite of client/server applications includes solutions for manufacturing, sales and distribution, financial accounting and human resource processes.

R/3 release 3.0 for AS/400 will provide a scalable solution for those customers wishing to manage their business processes on the AS/400.

Two benchmarks are primarily used today for making platform comparisons and sizing estimates of the R/3 product. The first is called FI and exercises the financial portion of the product. The second is called SD and exercises the sales and distribution portion of the product. Sizing information based on these benchmarks for R/3 release 3.0 is available from the IBM SAP Competency Centers.

For further details concerning a particular SAP configuration or for additional questions, please contact an IBM SAP Competency Center. The North America Competency Center can be reached at WASVMIC1(IBMSAPCC) or 1-800-426-0222. The International Competency Center can be reached at MUNIVM4(ISICC) or 49-6227-34-1298.

A paper describing current 2-tier and 3-tier benchmark results can be found at: <u>Http://www.sap.com/products/techno/media/pdf/wp_be2_e.pdf</u>

RPMark95

Information can be found at: http://www.cslinc.com/rpmark/rpmark.htm

NotesBench

Information can be found at: http://www.notesbench.org

19.2 Dynamic Priority Scheduling

On an AS/400 CISC-model, all ready-to-run OS/400 jobs and Licensed Internal Code (LIC) tasks are sequenced on the Task Dispatching Queue (TDQ) based on priority assigned at creation time. In addition, for N-way models, there is a cache affinity field used by Horizontal Licensed Internal Code (HLIC) to keep track of the processor on which the job was most recently active. A job is assigned to the processor for which it has cache affinity, unless that would result in a processor remaining idle or an excessive number of higher-priority jobs being skipped. The priority of jobs varies very little such that the resequencing for execution only affects jobs of the same initially assigned priority. This is referred to as Fixed Priority Scheduling.

For V3R6, the new algorithm being used is Dynamic Priority Scheduling. This new scheduler schedules jobs according to "delay costs" dynamically computed based on their time waiting in the TDQ as well as priority. The job priority may be adjusted if it exceeded its resource usage limit. The cache affinity field is no longer used in a N-way multiprocessor machine. Thus, on an N-way multiprocessor machine, a job will have equal affinity for all processors, based only on delay cost.

A new system value, QDYNPTYSCD, has been implemented to select the type of job dispatching. The job scheduler uses this system value to determine the algorithm for scheduling jobs running on the system. The default for this system value is to use Dynamic Priority Scheduling (set to '1'). This scheduling scheme allows the CPU resource to be spread to all jobs in the system.

The benefits of Dynamic Priority Scheduling are:

- No job or set of jobs will monopolize the CPU
- Low priority jobs, like batch, will have a chance to progress
- Jobs which use too much resource will be penalized by having their priority reduced
- Jobs response time/throughput will still behave much like fixed priority scheduling

By providing this type of scheduling, long running, batch-type interactive transactions, such as a query, will not run at priority 20 all the time. In addition, batch jobs will get some CPU resources rather than interactive jobs running at high CPU utilization and delivering response times that may be faster than required.

To use Fixed Priority Scheduling, the system value has to be set to '0'.

Delay Cost Terminology

• Delay Cost

Delay cost refers to how expensive it is to keep a job in the system. The longer a job spends in the system waiting for resources, the larger its delay cost. The higher the delay cost, the higher the priority. Just like the priority value, jobs of higher delay cost will be dispatched ahead of other jobs of relatively lower delay cost.

• Waiting Time

The waiting time is used to determine the delay cost of a job at a particular time. The waiting time of a job which affects the cost is the time the job has been waiting on the TDQ for execution.

• Delay Cost Curves

The end-user interface for setting job priorities has not changed. However, internally the priority of a job is mapped to a set of delay cost curves (see "Priority Mapping to Delay Cost Curves" below). The delay cost curve is used to determine a job's delay cost based on how long it has been waiting on the TDQ. This delay cost is then used to dynamically adjust the job's priority, and as a result, possibly the position of the job in the TDQ.

On a lightly loaded system, the jobs' cost will basically stay at their initial point. The jobs will not climb the curve. As the workload is increased, the jobs will start to climb their curves, but will have little, if any, effect on dispatching. When the workload gets around 80-90% CPU utilization, some of the jobs on lower slope curves (lower priority), begin to overtake jobs on higher slope curves which have only been on the dispatcher for a short time. This is when the Dynamic Priority Scheduler begins to benefit as it prevents starvation of the lower priority jobs. When the CPU utilization is at a point of saturation, the lower priority jobs are climbing quite a way up the curve and interacting with other curves all the time. This is when the Dynamic Priority Scheduler works the best.

Note that when a job begins to execute, its cost is constant at the value it had when it began executing. This allows other jobs on the same curve to eventually catch-up and get a slice of the CPU. Once the job has executed, it "slides" down the curve it is on, to the start of the curve.

Priority Mapping to Delay Cost Curves

The mapping scheme divides the 99 'user' job priorities into 2 categories:

• User priorities 0-9

This range of priorities is meant for critical jobs like system jobs. Jobs in this range will NOT be overtaken by user jobs of lower priorities. NOTE: You should generally not assign long-running, resource intensive jobs within this range of priorities.

• User priorities 10-99

This range of priorities is meant for jobs that will execute in the system with dynamic priorities. In other words, the dispatching priorities of jobs in this range will change depending on waiting time in the TDQ if the QDYNPTYSCD system value is set to '1'.

The priorities in this range are divided into groups:

- Priority 10-16
- Priority 17-22
- Priority 23-35
- Priority 36-46
- Priority 47-51
- Priority 52-89
- Priority 90-99

Jobs in the same group will have the same resource(CPU seconds and Disk I/O requests) usage limits. Internally, each group will be associated with one set of delay cost curves. This would give some preferential treatment to jobs of higher user priorities at low system utilization.

With this mapping scheme, and using the default priorities of 20 for interactive jobs and 50 for batch jobs, users will generally see that the relative performance for interactive jobs will be better than that of batch jobs, without CPU starvation.

Performance Testing Results

Following are the detailed results of two specific measurements to show the effects of the Dynamic Priority Scheduler:

In Table 19.1, the environment consists of the RAMP-C interactive workload running at approximately 70% CPU utilization with 120 workstations and a CPU intensive interactive job running at priority 20.

In Table 19.2 below, the environment consists of the RAMP-C interactive workload running at approximately 70% CPU utilization with 120 workstations and a CPU intensive batch job running at priority 50.

Table 19.1. Effect of Dynamic Priority Scheduling: Interactive Only						
	QDYNPTYSCD = '1' (ON)	QDYNPTYSCD = '0'				
Total CPU Utilization	93.9%	97.8%				
Interactive CPU Utilization	77.6%	82.2%				
RAMP-C Transations per Hour	60845	56951				
RAMP-C Average Response Time	0.32	0.75				
Priority 20 CPU Intensive Job CPU	21.9%	28.9%				

Table 19.2. Effect of Dynamic Priority Scheduling: Interactive and Batch						
	QDYNPTYSCD = '1' (ON)	QDYNPTYSCD = '0'				
Total CPU Utilization	89.7%	90.0%				
Interactive CPU Utilization	56.3%	57.2%				
RAMP-C Transations per Hour	61083	61692				
RAMP-C Average Response Time	0.30	0.21				
Batch Priority 50 Job CPU	15.0%	14.5%				
Batch Priority 50 Job Run Time	01:06:52	01:07:40				

Conclusions/Recommendations

- When you have many interactive jobs running on the system and want to ensure that no one CPU intensive interactive job 'takes over' (see Table 19.1 above), Dynamic Priority Scheduling will give you the desired result. In this case, the RAMP-C jobs have higher transaction rates and faster response times, and the priority 20 CPU intensive job consumes less CPU.
- Dynamic Priority Scheduling will ensure your batch jobs get some of the CPU resources without significantly impacting your interactive jobs (see Table 96). In this case, the RAMP-C workload gets less CPU utilization resulting in slightly lower transaction rates and slightly longer response times. However, the batch job gets more CPU utilization and consequently shorter run time.
- It is recommended that you run with Dynamic Priority Scheduling for optimum distribution of resources and overall system performance.

For additional information, refer to the Work Management Guide.

19.3 Main Storage Sizing Guidelines

To take full advantage of the performance of the new AS/400 Advanced Series using PowerPC technology, larger amounts of main storage are required. To account for this, the new models are provided with substantially more main storage included in their base configurations. In addition, since more memory is required when moving to RISC, memory prices have been reduced.

The increase in main storage requirements is basically due to two reasons:

- When moving to the PowerPC RISC architecture, the number of instructions to execute the same program as on CISC has increased. This does not mean the function takes longer to execute, but it does result in the function requiring more main storage. This obviously has more of an impact on smaller systems where fewer users are sharing the program.
- The main storage page size has increased from 512 bytes to 4096 bytes (4KB). The 4KB page size is needed to improve the efficiency of main storage management algorithms as main storage sizes increase dramatically. For example, 4GB of main storage will be available on AS/400 Advanced System model 530.

The impact of the 4KB page size on main storage utilization varies by workload. The impact of the 4KB page size is dependent on the the way data is processed. If data is being processed sequentially, the 4KB page size will have little impact on main storage utilization. However, if you are processing data randomly, the 4KB page size will most likely increase the main storage utilization.

The minimum memory available on RISC systems is 32MB versus 8MB on CISC systems. In most instances, 8MB CISC systems will require 32MB on RISC when running the same workload. However, if the 8MB CISC system is overcommitted in main storage utilization, then 64MB of main storage may be required on RISC.

As a first approximation of the main storage required when moving to AS/400 models using PowerPC AS technology, use the following guidelines. If you will be adding additional work as you upgrade to the new models, you should first determine what main storage would be required on CISC for this new workload before using the guidelines below.

Table 19.3. Main Storage Size Guidelines					
Main Storage Size on CISC	Main Storage Size on RISC				
Up to 160MB	(2 X CISC Main Storage) + 16MB (See note below)				
Greater than 160MB	2 X CISC Main Storage				
Note: The 16MB that is added is primarily due to the increase in size of the operating system code that must be resident in main storage. It is very important to take this increase into account when sizing systems with lower amounts of main storage.					

This initial estimate should be followed up by a more detailed analysis with BEST/1 for OS/400. Using BEST/1 to analyze your workload will take into account how main storage is being utilized on your current system.

Refer to WORKMEM for main storage sizing considerations for application development environments.

19.4 Memory Tuning Updates

The Performance Adjustment support (QPFRADJ system value) that is used for initially sizing memory pools and managing them dynamically at run time has been enhanced to support the new RISC hardware and V3R6. In addition, at V3R7 the CHGSHRPOOL and WRKSHRPOOL commands have been updated so that you can tailor memory tuning parameters used by QPFRADJ. Now you can specify your own faulting guidelines, storage pool priorities, and minimum/maximum size guidelines for each pool. This allows you the flexibility to set unique QPFRADJ parameters at the pool level.

The following changes were made for tuning done at IPL time (the system value QPFRADJ is set to 1 or 2):

- 1. The calculation for the minimum Machine pool size has been updated to reflect changes in the amount of storage needed for lines, controllers, and devices. The algorithm also has been changed to use (as a base value) the minimum Machine pool size calculated by the License Internal Code (LIC) instead of the tabular method (based on main storage size) that was used in previous releases.
- 2. The pool size calculation for the *INTERACT and *BASE pools has been updated. After the Machine and Spool pool sizes have been determined, 70% of the remaining storage is given to *INTERACT. The remaining 30% is given to the *BASE pool. To better support the unique demands of the client/server environment, on Server models the amounts are the opposite (70% to *BASE, 30% to *INTERACT).

The following changes were made to dynamic tuning done at run time - (the system value QPFRADJ is set to 2 or 3):

With V3R6, the page fault guidelines stated in the Work Management Guide (SC41-4306) are significantly higher than they are for V3R1. The large increases are due to now including index

faults in the count for V3R6, such as faults on storage management directories, user profiles, libraries, and file access paths.

- 1. The Dynamic Tuner no longer uses the guidelines published in the Work Management Guide for User pools (any pool except the Machine pool). The guidelines are on the average good, but the tuner is able to take more into account at run time to adjust the faulting guidelines based on workload characteristics. The Dynamic Tuner now calculates a run-time guideline based on the number of active jobs in the storage pool. The type of pool (*INTERACT, *SPOOL, or other shared pool) is also taken into consideration. The run-time guidelines could be much different than the published guidelines, especially if the number of jobs is much lower or much higher than the number of jobs for which the system is rated.
- 2. For Advanced Server models, the *BASE pool will be treated with higher priority than the *INTERACT pool. This means that if the Dynamic Tuner determines that both *BASE and INTERACT require more memory, *BASE will get more memory before *INTERACT. On traditional models, *INTERACT gets higher priority.
- 3. The minimum pool size for an active pool has been increased.
 - *INTERACT 3000K
 - *SPOOL 256K
 - *SHRPOOL1-10 1000K
 - If inactive, pools may be temporarily reduced to 256K (except the Machine and Base pool).
- 4. The maximum pool size for an active pool has been increased to 3072K for each active job in the pool (up from 2048K).

19.5 User Pool Faulting Guidelines

Due to the large range of AS/400 processors and due to an ever increasing variance in the complexity of user applications, paging guidelines for user pools are no longer published. Only machine pool guidelines and system wide guidelines (sum of faults in all the pools) are published. Even the system wide guidelines are just that...guidelines. Each customer needs to track response time, throughput, and cpu utilization against the paging rates to determine a reasonable paging rate.

There are two choices for tuning user pools:

- 1. Set system value QPFRADJ = 2 or 3. This algorithm has been changed for V3R6 and PTFed back to V3R1. The new algorithm is much better in several ways, including how it determines which pool has a paging problem, and the speed in which it can react to the needs of those pools with problems. Many customers are reporting they now don't have to worry about pool tuning at all. They still may need to analyze whether they need to increase total main storage. The rest of this section will help with that analysis.
- 2. Manual tuning. Move storage around until the response times and throughputs are acceptable. The rest of this section deals with how to determine these acceptable levels.

To determine a reasonable level of page faulting in user pools, determine how much the paging is affecting the interactive response time or batch throughput. These calculations will show the percentage of time spent doing page faults.

The following steps can be used: (all data can be gathered w/STRPFRMON and printed w/PRTSYSRPT). The following assumes interactive jobs are running in their own pool, and batch jobs are running in their own pool.

Interactive:

- 1. flts = sum of database and non-database faults per second during a meaningful sample interval for the interactive pool.
- 2. rt = interactive response time for that interval.
- 3. diskRt = average disk response time for that interval.
- 4. tp = interactive throughput for that interval in transactions per second. (transactions per hour/3600 seconds per hour)
- 5. fltRtTran = diskRt * flts / tp = average page faulting time per transaction.
- 6. flt% = fltRtTran / rt * 100 = percentage of response time due to
- 7. If flt% is less than 10% of the total response time, then there's not much potential benefit of adding storage to this interactive pool. But if flt% is 25% or more of the total response time, then adding storage to the interactive pool may be beneficial (see NOTE below).

Batch:

- 1. flts = sum of database and non-database faults per second during a meaningful sample interval for the batch pool.
- 2. flt% = flts * diskRt X 100 = percentage of time spent page faulting in the batch pool. If multiple batch jobs are running concurrently, you will need to divide flt% by the number of concurrently running batch jobs.
- 3. batchcpu% = batch cpu utilization for the sample interval. If higher priority jobs (other than the batch jobs in the pool you are analyzing) are consuming a high percentage of the processor time, then flt% will always be low. This means adding storage won't help much, but only because most of the batch time is spent waiting for the processor. To eliminate this factor, divide flt% by the sum of flt% and batchcpu%. That is: newflt% = flt% / (flt% + batchcpu%) This is the percentage of time the job is spent page faulting compared to the time it spends at the processor.
- 4. Again, the potential gain of adding storage to the pool needs to be evaluated. If flt% is less than 10%, then the potential gain is low. If flt% is greater than 25% then the potential gain is high enough to warrant moving main storage into this batch pool.

NOTE:

It is very difficult to predict the improvement of adding storage to a pool, even if the potential gain calculated above is high. There may be instances where adding storage may not improve anything because of the application design. For these circumstances, changes to the application design may be necessary.

Also, these calculations are of limited value for pools that have expert cache turned on. Expert cache can reduce I/Os given more main storage, but those I/Os may or may not be page faults.

19.6 Cryptography Performance

This section provides performance information for AS/400 Common Cryptographic Architecture Services/400 Version 3 (*). This information can be used to assist capacity planning for an AS/400. This data is not representative of a specific customer environment. Results in other environments may vary significantly.

This evaluation was completed on CISC models, but the relative performance and recommendations are similar.

Workload Description

CL and ILE/C application programs were used to utilize the cryptographic functions.

- Encipher and decipher a string of characters 100 times (DEA: Data Encryption Algorithm) with the key specified as:
 - ✤ key TOKEN
 - key label located in KEYSTORE
- CPB ("CPB and NetPerf Benchmark Descriptions" on page 0)
 - multiple send/receives pairs of 100 bytes
 - large data transfers

For information on CCAS/400 function, refer to the *Common Cryptographic Architecture Services*/400 *Installation and Operator's Guide*, SC41-0102.

Measurement Technique

The performance measurements were taken on a dedicated AS/400 (i.e., no other system load from other users) with an application program executing the primitive workload scenarios, previously described. For the communications environment, a second AS/400 was similarly configured. CPU utilization and run time were collected using the WRKSYSSTS command. CPU time is calculated (run time * CPU utilization = CPU time).

Measurement Results

Table 19.4. Cry	Table 19.4. Cryptographic Performance							
AS/400 Cryptographic Performance								
	Encipher and Decipher (100 times) From an Application							
		AS/400 D50, V3	R1, 2620 Cryptog	aphic Processor				
Data		Run		Delay		Operations		
Size	Key	Time	CPU Util	(ms per	CPU ms	per		
(Bytes)	Туре	(sec)	(%)	pair)	per pair	Second		
64	TOKEN	7	28.1	70	20	29		
1024	TOKEN	7	26.2	70	19	29		
2048	TOKEN	8	23.9	80	19	25		
4096	TOKEN	10	20.7	100	21	20		
8192	TOKEN	13	16.1	130	21	15		
16384	TOKEN	27	13.5	270	37	7		
64	KEYSTORE	22	78.4	220	173	9		
1024	KEYSTORE	28	79.6	280	223	7		
2048	KEYSTORE	28	77.1	280	216	7		
4096	KEYSTORE	30	73.3	300	220	7		
8196	KEYSTORE	34	65.8	340	224	6		
16384	KEYSTORE	48	49.8	480	239	4		
Note: 1 pair is 2	operations (1 enci	pher and 1 decipt	er)		•			

Table 19.5. SNA Session	Table 19.5. SNA Session Level Encryption					
	SNA Se	ssion Level Encryption	(SLE)			
	Арр	lication Program Trans	fer			
	CPB Benchma	rk with APPC and the I	CF Interface			
	AS/400 F25, V	V3R2, 2620 IOP, 2619 T	RLAN IOP			
		Data Rate	CPU Time	CPU Util		
Scenario	Time (sec)	(kbps)	(sec)	(%)		
10,000 100-byte snd/rcv	7					
SLE *ALL	1559	10	160	10		
SLE *NONE	245	65	94	38		
32M large transfer (send)						
SLE *ALL	259	1036	25	10		
SLE *NONE	28	9585	11	40		
Note: (1) Receives perfor	rm similarly to sends					

Conclusions/Explanations

- 1. The CPU issues requests to the cryptographic processor (CP). The CP then executes the cryptographic function and returns the results to the CPU. Because the CP handles these processing-intensive functions, the CPU is available to process other system activity.
- 2. The delay time and CPU time are roughly equivalent for both encipher and decipher functions. Therefore, the times listed for encipher and decipher pairs can be cut in half to estimate a single function (first table).
- 3. As the size of the encipher/decipher request (in bytes) increases, the delay time and CPU time increase (first table).
- 4. A key is presented to the CCAS/400 API as a 64-byte internal key token or as a key label of an internal key token which is located in a physical file, called key store.

- Encipher/decipher with a key specified as a TOKEN is faster and consumes less CPU time than when the key is specified as a label.
- Encipher/decipher with a key specified as a label located in KEYSTORE takes significantly longer and uses significantly more CPU time.
- 5. MAC (message authentication code) generate and MAC verify yield similar performance results to encipher and decipher. From lab measurements not provided here, trends from the different key specifications, absolute delay times, and CPU times are similar.
- 6. Note that the performance for the 16384-byte requests seem to have a higher than expected CPU time. This is because the maximum size that can be requested from the CPU to CP is 12K bytes. Therefore, requests that are larger than 12K bytes are broken up by the CCAS/400 API.
- 7. The CP can handle multiple cryptographic requests from different applications concurrently. From the 'operations per second' column in the table, note that the rates listed are for that single job (up to 29 operations per second. If multiple instances of the application were run, the aggregate rate would be higher. Being the performance tools do not provide the 2620 Cryptographic Processor utilization, it is difficult to project the capacity of the CP for particular customer scenarios. For small-sized operations, the IOP is limited to about 40 operations per second. For larger-sized operations, the IOP is limited to about 130K bytes per second.
- 8. When multiple cryptographic requests are sent to the CP from different jobs, they are processed serially in a first come, first served order. Note that if a long request arrives at the CP, it may delay requests from other jobs. This is particularly true with some of the PKA (Public-Key Algorithm) functions. If a job is performance sensitive, it can enter exclusive mode to prevent any other jobs from issuing cryptographic requests.
- 9. The 2628 Cryptographic Processor has similar performance characteristics to the 2620 CP. When processing requests, the 2628 CP uses the Commercial Data Masking Facility (CDMF) rather than the data encryption algorithm (DEA) that is used by the 2620 CP.
- 10. **SNA Session Level Encryption (SLE)** provides cryptography functions transparent to the application. By configuring a session for SLE (SLE parameter on the mode description), all transmissions for that session will be enciphered/deciphered. By using a CP, the CPU off-loads this CPU-intensive processing. There is, however, still a significant performance overhead for SLE (second table):
 - **Run Time**: Each time the application sends (or receives) data it must pass the data through the CP. Without SLE, a 100-byte send/receive pair takes 24.5 ms; with SLE, it takes 155.9 ms (6 times longer). Being each send/receive pair transaction uses the CP four times (twice on the system under test, twice on the other AS/400), each use of the CP adds about 33 ms of additional delay. This seems like a large percentage increase here, as this application does no processing other than sending data. For real applications that are more complex, this extra delay may be less significant.

The delay times will vary based on the size and type of the request to the CP and the CPW (relative system performance metric) of the CPU.

- **CPU Time**: When SLE is used, there is also additional processing on the CPU to facilitate getting the data to and from the CP. For the send/receive scenario, SLE used 1.7 times more CPU. For the large transfer scenario, SLE used 2.3 times more CPU. Again, depending on the application that is associated with the transmissions, this overhead may be significant or unnoticeable.
- **Data transfer rate**: For large transfers, the impact of using SLE is significant. For the case in the second table, SLE slowed the data rate by almost 10-fold. For LAN environments (very efficient, and high speed), the effect of SLE is significant. For WAN environments with linespeeds of 64 kbps or slower, the slow-down effect of SLE may be insignificant. Actual data rates will vary based on the size and type of the transmission.
- 11. There are trade-offs to be considered with encryption over communications. For best performance, use a limited amount of encryption from the application (with the CP) and then send the data. For maximum encryption and ease-of-use, use SLE. If most of the data is to be encrypted, use SLE; SLE provides better performance than application (with the CP) encryption in this case. For a send/receive pair, note from the first table that application encryption uses 20 ms of CPU time. Note from the second table that SLE consumes only 16 ms of CPU time per send/receive pair.

19.7 AS/400 NetFinity Capacity Planning

Performance information for AS/400 NetFinity attached to a V4R1 AS/400 is included below. The following NetFinity functions are included:

- Time to collect software inventory from client PCs
- Time to collect hardware inventoy from client PCs

The figures below illustrate the time it takes to collect software and hardware inventory from various numbers of client PCs. This test was conducted using the Rochester development site, during normal working hours with normal activity (ie. not a dedicated environment). This environment consists of:

- 16 and 4Mb token ring LANs (mostly 16)
- LANs connected via routers and gateways
- Dedicated AS/400
- TCP/IP
- Client PCs varied from 386s to Pentiums (mostly 100 MHz with 32MB memory), using OS/2, Windows/95 and NT
- About 20K of data was collected, hardware and software, for each client

While these tests were conducted in a typical work environment, results from other environments may vary sigificantly from what is provided here.


Figure 19.1. AS/400 NetFinity Software Inventory Performance



Figure 19.2. AS/400 NetFinity Hardware Inventory Performance

Conclusions/Recommendations for NetFinity

- 1. The time to collect hardware or software information for a number of clients is fairly linear.
- 2. The size of the AS/400 CPU is not a limitation. Data collection is performed at a batch priority. CPU utilization can spike quite high (ex. 80%) when data is arriving, but in general is quite low (ex. 10%).
- 3. The LAN type (4 or 16Mb Token Ring or Ethernet) is not a limitation. Hardware collection tends to be more chatty on the LAN than software collection, depending on the hardware features.
- 4. The communications protocol (IPX, TCP/IP, or SNA) is not a limitation.
- 5. Collected data is automatically stored in a standard DB/2/400 database file, accessible by SQL and other APIs.
- 6. Collection time depends on clients being powered-on and the needed software turned on. The server will retry 5 times.
- 7. The number of jobs on the server increases during collection and decreases when not needed.

Chapter 20. General Performance Tips and Techniques

For this version of the Performance Capabilities Guide, there is only a single entry in this chapter. Over time, this section should grow to cover a variety of useful topics that "don't fit" in the document as a whole, but provide useful things that customers might do or deal with special problems customers might run into on the AS/400.

20 .1 Adjusting Your Performance Tuning for Threads

History

Historically, the AS/400 programmer has not had to worry very much about threads. True, they were introduced into the machine some time ago, but the average RPG application does not use them and perhaps never will, even if it is now allowed. Multiple-thread jobs have been fairly rare. That means that those who set up and organize AS/400 subsystems (e.g. QBATCH, QINTER, MYOWNSUBSYSTEM, etc.) have not had to think much about the distinction between a "job" and a "thread."

The Coming Change

But, threads are a good thing and so applications are increasingly using them. Especially for customers deploying (say) a significant new Java application, or Domino, a machine with the typical one-thread-per-job model may suddenly have dozens or even hundreds of threads in a particular job. Unfortunately, they are distinct ideas and certain AS/400 commands carefully distinguish them. If AS/400 System Administrators are careless about these distinctions, as it is so easy to do today, poor performance can result as the system moves on to new applications such as Lotus Domino or especially Java.

With Java generally, and with certain applications, it will be commonplace to have multiple threads in a job. That means taking a closer look at some old friends: MAXACT and MAXJOB.

Recall that every subsystem has at least one pool entry. Recall further that, in the subsystem description itself, the pool number is an arbitrary number. What is more important is that the arbitrary number maps to a particular, real storage pool (*BASE, *SHRPOOL1, etc.). When a subsystem is actually started, the actual storage pool (*SHRPOOL1), if someone else isn't already using it, comes to life and obtains its storage.

However, storage pools are about more than storage. They are also about job and thread control. Each pool has an associated value called MAXACT that also comes into play. No matter how many subsystems share the pool, MAXACT limits the total number of threads able to reside and execute in the pool. Note that this is *threads* and not *jobs*.

Each subsystem, also, has a MAXJOBS value associated with it. If you reach that value, you are not supposed to be able to start any more jobs in the subsystem. Note that this is a *jobs* value and not a *threads* value. Further, within the subsystem, there are usually one or more JOBQs in the subsystem. Within each entry you can also control the number of jobs using a parameter. Due to an unfortunate turn in history, this parameter, which might more logically be called MAXJOBS today is called MAXACT. However, it controls *jobs*, not *threads*.

Problem

It is too easy to use the overall pool's value of MAXACT as a surrogate for controlling the number of Jobs. That is, you can forget the distinction between jobs and threads and use MAXACT to control the activity in a storage pool. But, you are not controlling jobs; you are controlling threads.

It is also too easy to have your existing MAXACT set too low if your existing QBATCH subsystem suddenly sees lots of new Java threads from new Java applications.

If you make this mistake (and it is easy to do), you'll see several possible symptoms:

- **V** Mysterious failures in Java. If you set the value of MAXACT really low, certainly as low as one, sometimes Java won't run, but it also won't always give a graceful message explaining why.
- V Mysterious "hangs" and slowdowns in the system. If you don't set the value pathologically low, but still too low, the system will function. But it will also dutifully "kick out" threads to a limbo known as "ineligible" because that's what MAXACT tells it to do. When MAXACT is too low, the result is useless wait states and a lot of system churn. In severe cases, it may be impossible to "load up" a CPU to a high utilization and/or response times will substantially increase.
- ▼ Note carefully that this can happen as a result of an upgrade. If you have just purchased a new machine and it runs slower instead of faster, it may be because you're using "yesterday's" limits for MAXACT.

Solution

Make sure the storage pool's MAXACT is set high enough for each individual storage pool (CHGSHRPOOL command or WRKSYSSTS command to fix it). A MAXACT of *NOMAX will sometimes work quite well, especially if you use MAXJOBS to control the amount of working coming into each subsystem.

Use MAXJOB in the subsystem to control the amount of outstanding work:

CHGSBSD QBATCH MAXJOBS(newmax)

Use the Job Queue Entry in the subsystem to have even finer control of the number of jobs:

CHGJOBQE SBSD(QBATCH) JOBQ(QBATCH) MAXACT(newqueue job maximum)

Note in this particular case that MAXACT does refer to jobs and not threads.

Appendix A. CPW Benchmark Description

CPW is designed to evaluate a computer system and associated software in the commercial environment. It is rigidly defined for function, performance metrics, and price/performance metrics. It is NOT representative of any specific environment, but it is generally applicable to the commercial computing environment.

- What CPW is
 - Test of a range of data base applications, including simple and medium complexity updates, simple and medium complexity inquiries, realistic user interfaces, and a combination of interactive and batch activities.
 - Test of commitment control
 - Test of concurrent data access by large numbers of users running a single group of programs.
 - Reasonable approximation of a steady-state, data base oriented commercial application.
- What CPW is not:
 - ✤ An indication of the performance capabilities of a system for any specific customer situation
 - ✤ A test of "ad-hoc" (query) data base performance
- When to use CPW data
 - Approximate product positioning between different AS/400 models where the primary application is expected to be oriented to traditional commercial business uses (order entry, payroll, billing, etc.) using commitment control

CPW Application Description

There are five business functions of varying complexity that are simulated in the CPW application. Each simulated user has access to all five functions and must exercise all five with a weighted-random sequence during the course of a measurement. A "business transaction" is defined as a combination of a menu selection transaction and one of the five transactions described below.

1. New Order Transaction

The New-Order transaction enters a customer order for 5 to 15 items from a set of supply warehouses. It is a medium-weight read-write transaction that is the foundation of the CPW benchmark. The transaction is executed approximately 43% of the time from each terminal. There is a restriction that 90% of the transactions complete within a 5-second end-user response time.

The transaction accesses most of the files in the data base with a combination of reads, read-with-updates, and insertion of new records

Some of the transactions are required to fail, measuring the effect of transaction recovery on the overall system performance.

The New-Order transaction uses a display screen that uses a majority of a 24x80 character display.

2. The Payment Transaction

Like the New-Order transaction, the Payment transaction is executed approximately 43% of the time.

The basic transaction accesses 3 files by retrieving and updating a single record in each file and inserts a record into a fourth file. 60% of the time, additional reads of a file are required to determine the appropriate key values for the accesses of the basic transaction. This simulates the selection of a customer name from a list of names that match input from an end-user.

The Payment transaction also uses a full-screen interface, although it is not as complex as the New-Order transaction's screen. Most of the Payment display is used for presenting output of information that is read from the files.

3. The Order Status Transaction

The Order-Status transaction is a medium-weight read-only transaction that has a relatively low frequency of execution (4%).

The input on the display is quite simple, with the bulk of the transaction spent reading from 3 files and formatting a full screen of output information. Like the payment transaction, some of the time the input requires arbitrary selection of a "customer" from a list (requiring input of multiple customers with duplicate last names) and some of the time the input will specify an exact set of key-values to be used to retrieve the order-status information.

4. The Delivery Transaction

The Delivery transaction has two parts: a trivial interactive request that identifies a warehouse and carrier to be used to make deliveries and a more significant batch operation that identifies "deliverable" orders assigned to that warehouse and processes a set of 10 of those orders. It is a medium-weight read-write transaction that has a relatively low frequency of execution (4%). The transaction is intended to represent a longer-running (under 80 seconds) operation that would not normally be done at an interactive display, but would often be done at the same time as higher-priority interactive tasks are done and would operate against the same data base as the interactive transactions.

The Delivery transaction accesses six of the data base files with a combination of read, read-update, and delete. It also posts information in a log file (insert), primarily for the purpose of maintaining statistics for validation of the benchmark.

5. The Stock Level Transaction

The Stock-level transaction is a heavy-weight read-only transaction that has a low frequency of execution (4%). It has very simple input and output to the display screen, but it includes up to 400 read operations across 3 files and it generates a count of entries that satisfy some specific criteria. While not as complex as a true inventory-control-type application, it exercises the same type of file I/O and logic that would be associated with such an application.

CPW Data Base

There are nine data base files used in the CPW benchmark. The files are widely varied in size and are required to scale with the amount of throughput that is claimed for a given system. Some of the files in CPW experience updates to their indexes as well as to the content of their data. This is done through the use of inserts and deletions to the files during the course of the benchmark. All fields in the files contain information that is used by some portion of the benchmark.

1. Item File

- 50,000 records constant regardless of size of rest of data base
- 1 key field
- 4 fields total
- Records of 35-60 bytes
- Read by New-order

2 Warehouse File

- 1 record per warehouse (10 terminals are required to support each warehouse)
- 1 key field
- 9 fields total
- Records of 50-80 bytes
- Read by New-order, updated by Payment

3 **District File**

- 10 records per warehouse (1/terminal)
- 2 key fields
- 11 fields total
- Records of 70-100 bytes
- Read by Stock-level, updated by New-order and Payment

4. **Customer File**

- 30,000 records per warehouse
- 3 key fields
- 21 fields total
- Records of 240-370 bytes
- Read by New-order, Payment, and Order-status, and updated by Payment and Delivery

5. Stock File

- 50,000 records per warehouse
- 2 key fields
- 17 fields total
- Records of 300-320 bytes
- Read by Stock-level, updated by New-order

6. History File

- Starts at 30,000 records per warehouse
- Each Payment transaction adds one record
- 3 key fields
- 6 fields total
- Records of 35-50 bytes
- Added to by Payment

7. Order File

- Starts with 30,000 records per warehouse
- Each New-order transaction adds one record
- 4 key fields
- 8 fields total
- Records of about 30 bytes
- Added to by New-order, read by Order-status, updated by Delivery

8. Order-Line File

- Starts with 300,000 records per warehouse
- Each New-order transaction adds an average of 10 records
- 4 key fields
- 10 fields total
- Records of about 50 bytes
- Added to by New-order, read by Order-status and Stock-level, updated by Delivery

9. **New-Order File**

- Starts with 2,000 records per warehouse
- Continuously added to and deleted from
- 3 key fields
- 3 total fields
- Records of about 10 bytes
- Added to by New-order, deleted from by Delivery

CPW Benchmark Summary

Transaction		Files		
Туре	Complexity	Accessed	Logical	Activity
New Order	Complex	Warehouse	1	Get warehouse tax
		Customer	1	Get customer discount
		District	2	Get & update next order ID
		Order	1	Add a record
		New-Order	1	Add a record
		Item	10	Find cost of 10 items
		Stock	20	Adjust stock qty and YTD ordered
		Order-Line	10	Add 10 records
Payment	Simple	Warehouse	2	Get & update warehouse YTD
		District	2	Get & update district YTD
		Customer	2	Get & update YTD balance & count
		History	1	Add one record
Order Status	Medium	Customer	1	Get balance
		Order	1	Get order list
		Order-Line	10	Get all order lines
Delivery	Complex	New-Order	20	Get & delete 10 records
		Order	20	Get/update 10 records with carrier ID
		Order-Line	200	Get & update 100 records with date
		Customer	20	Get/update balance & delivery count
Stock-Level	`Medium	District	1	Get next order ID
	to	Order-Line	200	Get lines from last 20 orders
	Complex	Stock	200	Get quantity on hand from 200 items

This table summarizes activities performed within the CPW benchmark:

<u>CPW Required Functions</u>

• ACID Properties

Atomicity, Consistency, Isolation, and Durability properties are defined as the minimal integrity required of any system for which benchmark results are to be published. Essentially, these properties are designed to ensure that the system under test can guarantee that transaction integrity will be maintained (either all or none of a transaction is committed to the data base) through normal processing and in the event of a single point of failure (such as a power failure or disk failure).

• Configuration Size Requirements

For each warehouse configured in the data base, all the other files (except the constant Item file) must be scaled according to the descriptions listed above. Also, 10 workstations must be configured on the system and must exercise the applications during the benchmark period.

For the claimed throughput rating of the system measured, sufficient disk space must be configured to hold 180 days of historical data (from the continuously expanding History, Order, and Order-line files) and its associated indexes. There must also be sufficient space in the journal log files to run at the claimed throughput for an 8-hour period.

• Other Requirements

The benchmark specification allows flexibility in implementation, as long as the functions described in

the transaction descriptions are completed. Each transaction has a specific set of actions that must be taken using data from the input screen and from the data base files. The means of achieving those actions is left largely to the implementer.

The design of each workstation screen is more rigidly specified than the application logic. Each screen is described in detail within the specification. The screen interface is required to support such functions as moving the cursor from field to field, correcting data that has already been keyed, requiring input in some fields, and other things that are associated with a full-screen interface.

The keyed entry for each transaction is also rigidly specified. For any given transaction, specific rules are defined for the rate of keying input, the type and value (often a random number or weighted-random number) of the input, and the amount to delay after a response is received before selecting the next transaction.

Transaction selection is also rigidly specified. At least 43% of the transactions must be Payment transactions and at least 4% each of the transactions must be Order-status, Delivery, and Stock-level transactions. The remaining 45% of the transactions may be a mixture of New-orders and the other four transactions, although it is to the advantage of the measurement to have as many New-orders as possible.

CPW Metrics

The performance metric of CPW is the relative system throughput for the New-order transaction. This is expressed as the Relative System Performance Metric. Since the other four transactions have minimum frequency requirements and since the New-order transaction is the primary transaction in the benchmark, the throughput is measured in terms of New-orders only. A valid CPW Value is achieved when the correct transaction mix is used and at least 90% of the New-order transactions complete within 5 seconds.

Appendix B. AS/400 Sizing

In this section three sizing tools are discussed.

This section covers:

• IBM Workload Estimator for AS/400 (Available 8/03/99)

The Workload Estimator will estimate the proper sized AS/400 for Domino, Java, Net.Commerce and traditional workloads, individually or in combination.

• the AS/400 Capacity Planner (BEST/1 for the AS/400)

Best for MES upgrade sizing, or complex 'new business' system sizing.

• the AS/400 BATCH400 tool,

Best for MES upgrade sizing where the 'Batch Window' is important.

B.1 IBM Workload Estimator for AS/400

The new Estimator for AS/400 is available for sizing Domino and other workloads. To access the Estimator do the following:

For IBM employees:

1. Start at the Server Sales web page http://w3.ibm.com/server/sales

2.Select a region plus select the AS/400 server

3. Select Proposal Resources...

4. Expand the Tools download section...

5. Choose the "AS/400 Workload Estimator and Sizing Resources" page.

6. Web executable and downloadable forms of the Workload Estimator tool are available at the start of the page.

For IBM Business Partners

1. Start at the PartnerInfo Web page: <u>http://partners.boulder.ibm.com</u>

2. Select Shortcuts...

- 3.Select All IBM Servers (Server Sales)...
- 4.Select the AS/400 tab...
- 5. Select Proposal Resources...
- 6. Expand the Tools download section...
- 7. Choose the "AS/400 Workload Estimator and Sizing Resources" page.
- 8. Web executable and downloadable forms of the Workload Estimator tool are available at the start of the page.

Customers should contact their IBM or Business Partner for assistance in accessing and using the Workload Estimator tool.

B.2 AS/400 CAPACITY PLANNER (BEST/1 for the AS/400)

BEST/1 for the AS/400 BEST/1 for the AS/400 is the product of an alliance with BGS Systems, Inc, and will continue to be a part of the IBM Performance Tools/400 product. The capacity planner gives predicted performance information for response times, throughputs, and device utilizations based on estimated and/or measured workloads with a system configuration.

What It Does

The capacity planner helps to analyze the present and future performance requirements for the AS/400 system. The capacity planner allows the use of predefined profiles and/or measured data to create an environment similar to the application environment required.

Use the predefined profiles for an initial proposal. Use the measurement capability alone if the current activity is growing or being analyzed. Mix the predefined profiles with the measured data if new applications are being added or the current ones are being changed significantly. The workloads are then mixed based on the number of local and/or remote devices specified. Optionally, the user can specify a response time or throughput objective for each of the workloads. These objectives (maximum for response time and minimum for throughput) represent the performance requirements.

After the workload has been defined, the capacity planner uses the measured configuration or allows the user to select from an IBM supplied list of configurations. The configuration and workload are analyzed and modeled to predict performance parameters such as response times, throughputs, and device utilizations. When measured configurations are not available, BEST/1 models perspective hardware configurations based on service times measured from a RAMP-C environment.

The capacity planner's evaluator then compares these numbers against a set of utilization guidelines and the optional response times or throughput objectives. If either the guidelines or objectives are not met, the evaluator recommends an upgrade to the system and reevaluates the adjusted system. This iterative process continues until a configuration is found that satisfies the guidelines and objectives.

Additionally, the planner includes a system growth function. The growth function allows the user to specify an anticipated growth rate over the entire system or by specific workloads. The capacity planner then estimates what configuration changes are required to sustain performance over time.

What Is Supported

The actual AS/400 workload can be measured using the AS/400 Performance Monitor. BEST/1 uses this data to model system activities and provide workload support for the normal AS/400 environment, and also functions such as PC Support (Work Station Feature and Shared Folders) and Display Station Passthrough (Source and Target).

BEST/1 also includes a set of predefined workloads which can be used to represent applications and workloads which are not measured. The predefined profiles include RAMP-C, Officevision/400, RTW, Batch, and Spool.

Support is provided for the various system functions, including checksum protection, purge option, and disk mirroring. In addition, BEST/1 also supports multiple memory pools, multiple priorities, multiple

ASPs, batch, batch and interactive relationships, and the ability to model hardware enhancements the day they are announced.

BEST/1 for the AS/400 allows batch job analysis to be based on pool, priority, pathlength, and I/O characteristics the same as can be done for interactive jobs. This allows the user to set objectives for batch throughput, independent of interactive work, or in relation to interactive work.

BEST/1 provides a rated throughput for batch expressed in transactions per hour. This information can be used to estimate changes in throughput based on configuration or workload changes. BEST/1 does not provide detailed batch window analysis or job scheduling analysis. For modeling help in this area, reference B.2, "BATCH400" on page 274.

The capacity planner can also be used to assist the System/36 customer in selecting an appropriately sized IBM AS/400 system to meet their performance requirements. The capacity planner works with a System/36 migration utility procedure which is part of the System/36 release 5.1 coexistence PTF package (DK3700 or later) and System/36 Release 6. The utility uses the S/36 measured performance data created by the System Measurement Facility (SMF), and through modeling, translates the data into AS/400 System/36 environment performance data for the AS/400 Capacity Planner. The capacity planner can then be used to determine an AS/400 configuration that meets the anticipated performance needs.

Where to Get It

This capacity planner is part of the previously mentioned IBM AS/400 Performance Tools package which is a licensed program for the AS/400 system (5763-PT1 in V3R1 and 5716-PT1 in V3R6). This package also includes the measurement facilities needed to use the measurement interface capabilities of the capacity planner. This product's users guide includes more details on the measurements and capacity planner as well as specifics for the Capacity Planner System/36 Migration Utility option (*IBM AS/400 Programming: Performance Tools Guide*, SC41-8084 and *IBM AS/400 BEST/1 Capacity Planning Tool Guide*, SC41-3341). Other references also include the Capacity Planning Redbook (GG24-3908), and the BEST/1 Educational Video package (SK2T-6740-00 for 1/2" open reel, SK2T-6741-00 for 1/2" cartridge, SK2T-6742-00 for 1/4" cartridge, and SK2T-6743-00 for 8mm data tapes).

A Skill Dynamics Education course is available. The course is number CEM4930C - "AS/400 Capacity Planning using BEST/1."

Announcement PTF summaries for BEST/1

For the latest PTFs that are available to support the AS/400 Advanced Series with PowerPC AS technology, refer to HONE item RTA000089352. If you do not have access to HONE, please contact an IBM representative for this information.

V3R6 Enhancements and PTF Changes for V3R1, V3R0.5, and V2R3

• Analysis Enhancements

With the October Announcement, BEST/1 now supports the entire range of PowerPC-based Advanced Series AS/400 models.

Specifically this includes the following systems:

510 - 2144, 2143 500 - 2142, 2141, 2140 400 - 2133, 2132, 2131, 2130 50S - 2121, 2120 40S - 2110 530 - 2153, 2152, 2151, 2150 53S - 2156, 2155, 2154

For V3R1, two new "family" keywords have been added to the CPU Model Definition:

- ✤ *POWERAS 510, 500, and 400 models
- ✤ *POWERSRV 50S and 40S models

In V3R1, the "Upgrade to family" keyword for all CISC-based processors has been intentionally left as *ADVSYS or *ADVSRV. Explicit action is required by the user to configure BEST/1 to automatically trigger upgrade recommendations to PowerPC AS systems by changing the "Upgrade to family" keyword to *POWERAS or *POWERSRV on the processors where this is desired.

In V2R3, there is no "Upgrade to family" setting. BEST/1 only upgrades to CPU models which are currently available. As long as Advanced Series and Advanced Servers have their "Currently available" keyword set to Y=Yes, other CPU models will upgrade to them, because they appear earlier in the Hardware characteristics menu than the PowerPC AS models (with the exception of the 30S 2411 and 2412 models). If one wants BEST/1 to automatically trigger upgrade recommendations to PowerPC AS systems, one must change the "Currently available" keyword of the Advanced Series and/or Advanced Server models to N=No and change the PowerPC AS models to Y=Yes. V2R3 also contains internal rules which allow Server models to only upgrade to and from other Server models. Portables can not be upgraded to or from other CPU models.

BEST/1 now provides the capability to model the transition between CISC-based (Complex Instruction Set Computer) and RISC-based (Reduced Instruction Set Computer) or PowerPC technology based systems. IBM supplied values (not currently user-modifiable) are provided to manage the conversion at a transaction-level for the following values:

- ✤ CPU time
- ✤ Working Set Size
- ✤ I/O counts

These conversion factors are specified for General purpose, CPU intensive, I/O intensive, and Development workloads and are applied during analysis when BEST/1 models are analyzed with a configuration containing a RISC-based CPU model. The current conversion factors are listed in Table B.1 below:

Table B.1. CISC-to-RISC Conversion Factors

		Working Set	
Description	CPU Time	Size*	I/O Counts
General purpose workload type	1.00	2.00	1.00
CPU intensive workloads where the CPU per I/O rate is high	0.50	1.50	1.00
I/O intensive workloads where the CPU per I/O rate is low	1.00	2.30	1.00
Development workload where compile and debug are the	1.50	5.00	1.00
bulk of the work			

Note:

(*) In addition to the working set size conversion factor, the minimum RISC machine pool requirement of 16 MB (regardless of total system mainstore memory size) is included when recommending pool sizes during a CISC-to-RISC upgrade. In extreme cases this will cause excess memory to be recommended during a growth analysis because BEST/1 preserves the ratio of the machine pool to total mainstore memory as the other pools are increased in size due to workload growth.

DISCLAIMER: CONVERSION FACTORS AND BEST/1 RESULTS HAVE NOT BEEN SUBMITTED TO ANY FORMAL IBM TEST AND ARE DISTRIBUTED ON AN "AS IS" BASIS AT THIS TIME WITHOUT ANY WARRANTY EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PARTICULAR PURPOSE. The use of BEST/1 results is a customer responsibility and is dependent on the customer's operational environment; customers applying BEST/1 results do so at their own risk. Conversion factors may change as additional experience is gained by analyzing comparable performance data on CISC and RISC systems. New PTFs will be made available as appropriate.

*NORMAL and *BATCHJOB are the default and recommended settings for workload type and, furthermore, are the only workload types available in V2R3 and V3R0.5. These types will refer to conversion factors which are expected to provide good results on most general purpose customer configurations. The *TRNxxxx workload types represent specific workload characteristics, and are provided for users who understand how they match the workloads in their model per Table B.2 below:

Table B.2. Workload Type Characteristics		
Description	Normal	Batch
General purpose workload type	*NORMAL	*BATCHJOB
CPU intensive workloads where the CPU per I/O rate is high	*TRNNORM1	*TRNBAT1
I/O intensive workloads where the CPU per I/O rate is low	*TRNNORM2	*TRNBAT2
Development workload where compile and debug are the bulk of the	*TRNNORM3	*TRNBAT3
work		

Additional description:

*NORMAL and *BATCHJOB

These default workloads are represented by traditional commercial applications. CPU profiles for these workloads have up to 10-20% of their CPU time spent in application programs and the remaining 80+% spent in operating system programs. This is because typical RPG and COBOL business applications utilize a significant amount of system services such as database I/O, query processing, workstation/printer processing, and communication I/O.

*TRNNORM1 and *TRNBAT1

These types of workloads are referred to as Application Compute Intensive. In these workloads, a majority of the CPU time is is spent in application programs. Although these applications can be written in other languages, they are typically written in in ILE C. Examples of these types of applications are: financial modeling applications which do a significant amount of numeric calculations, statistical analysis applications, 4GL interpreters, and applications which implement complex business rules.

✤ *TRNNORM2 and *TRNBAT2

These types of workloads are characterized on a properly configured and tuned system by the majority of application time being spent doing I/O.

*TRNNORM3 and *TRNBAT3

These types of workloads are characterized by the OPM development environment where compiles are done with *NOPTIMIZE. Specifying *OPTIMIZE can increase the CPU time. For ILE development environments use the *NORMAL workload type.

To aid in the classification of a workload Table B.3 below illustrates some sample CPU per I/O values for various DASD response times. One can analyze their BEST/1 workload details to determine if any workload types should be changed. Calculate the CPU per I/O value by dividing the CPU seconds per transaction by the total reads and writes per transaction. For example, if the workload has a single function executing 1 function per user, the function defines a single transaction executing 1 transaction per function, and the transaction specifies 100 I/Os and 5 CPU seconds, then it is characterized by a CPU per I/O value of 0.05 (5 / 100 = 0.05). Thus if the average DASD response time for the configuration is 20 milliseconds this would indicate the *NORMAL classification is the correct setting for the workload.

(Note: if there is more than 1 transaction definition in the workload, one will need to account for the number of transactions in the calculation such that the CPU per I/O value is a weighted average: if 10

transactions have a value of 0.05 and 40 transactions have a value of 0.15, then the weighted average value is 0.13).

Table B.3. Sample CPU per I/O Values and the Appropriate Workload Type						
	I/O Intensive *TRNNORM2	*NORMAL	CPU Intensive *TRNNORM1			
	or	or	or			
Avg I/O Resp Time	*TRNBAT2	*BATCHJOB	*TRNBAT1			
10 msec	< 0.003	0.003 to 0.040	> 0.040			
20 msec	< 0.005	0.005 to 0.080	> 0.080			
30 msec	< 0.008	0.008 to 0.120	> 0.120			

• Performance Enhancements for V3R6

BEST/1 now runs in ILE instead of EPM, resulting in better performance.

• Model Creation Enhancements for V3R6

BEST/1 now supports the ADV36 Job Type.

• Usability Enhancements for V3R6

BEST/1 supports changing the feature of multiple IOP's, controllers and arms during a single operation.

• Predefined Workloads for V3R6

Additional server predefined workloads have been added to provide workloads for multimedia environments.

B.3 BATCH400

BATCH400 is a tool for Batch Window Analysis available for V3R6+ systems. It is an internal use only tool at this time. Instructions for requesting a copy are at the end of this description.

BATCH400 is a tool to enable AS/400 batch window analysis to be done using information collected by the OS/400 performance monitor.

BATCH400 addresses the often asked question: 'What can I do to my system in order to meet my overnight batch run-time requirements (also known as the Batch Window).'

The BATCH400 tool creates a 'model' from AS/400 performance data. This model will reside in a file called 'QSBSCHED' in the target library. The tool can then be asked to analyze the model and provide results for various 'what-if' conditions. Individual batch job run-time, and overall batch window run-times will be reported by this tool.

BATCH400 Output description:

- 1. Configuration summary shows the current and modeled hardware for DASD and CPU.
- 2. Job Statistics show the modeled result followed by the original (probably measured) data for each workload. Workloads are given short names (like b6) that represent either a single batch job or a collection of jobs grouped together. A listing at the end of the output shows the job number/user/name associated with each workload name. A short name like b6 indicates that a job is in the 6th 'thread' of jobs, since the letter is 'b', it is the second job in the thread (a6 being the first). Most other fields in this section are self explanatory. (Tr/Sec) is the Sync I/O per second rate for batch jobs. (Tr/Sec) is the Interactive transaction/second for the interactive workload. (Int Rt) is the Interactive response time. (Bat eff) is the batch efficiency of the workload, a value of 1.0 means the the workload is 100% CPU bound. (ExWait) is the time we had to add to the workload to account for the entire time the job was present in the AS/400, large values here can indicate either delay jobs, or excessive DB contention.
- 3. Thread summary shows the start/stop/elapsed time for entire threads.
- 4. Graph of Threads vs. Time of Day shows a 'horizontal' view of all threads in the model. This output is very handy in showing the relationship of job transitions within threads. It might indicate opportunities to break threads up to allow jobs to start earlier and run in parallel with jobs currently running in a sequential order.
- 5. Total CPU utilization shows a 'horizontal' view of how busy the CPU is. This report is on the same time-line as the previous Threads report.
- 6. Bar chart of Thread elapsed Times shows a comparison of all threads based on end-to-end thread run time. All threads start at time of zero and reach up into this report depending on how long they run. This report shows a 'vertical' depiction of each thread, and will occasionally show better job transition details than the 'horizontal' view noted above. This report can help to identify the longest running thread and possible candidate workloads to be improved, or scheduled for different times.
- 7. The rest of the output shows the model that is stored in QSBSCHED. This is the model that was used for the analysis. The config summary, and workload details are followed by a listing of the workload definitions. This workload definition usually shows:
 - interactive workloads which are a summary of all interactive jobs at a given priority level (type 1 workloads).
 - System workloads which are a summary of various VMC tasks and 'short running' batch jobs at a given priority level (type 3).
 - Batch workloads which are a summary of individual batch jobs which are usually the focus of the analysis (type 2).
 - Async workloads which are a summary of all the Asynchronous I/O tasks running on the system (type 4).

After looking at the results, use the BATCH400 *CHANGE option to invoke WRKOBJPDM. This will allow you to edit the model. You can alter the job dependencies. Maybe job b6 doesn't have to follow job

a6 in thread 6. You can remove the previous job linkage for b6 and rename b6 job to b12 (where 12 is the next available thread number). Upon saving and exiting, BATCH400 will analyze the new model.

For now this tool is an internal use only tool that can be requested by issuing the following command from a VM command line:

REQUEST BATCH400 FROM DENNEY AT RCHVMV2

Appendix C. DASD IOP Device Characteristics

This appendix describes the DASD models supported by the 6502, 6512,6530, 6532, 6751 and 6754 IOPs and the 2726, 2740, 2741, 2748 and 9728 IOAs.

6502/6512 Disk Unit Controller for RAID

Feature #6502 is a disk controller with a 2MB write-cache. Feature #6512 is an enhanced disk controller with a 4MB write-cache. Both can provide RAID-5 protection for up to 16 internal disk units installed in the Storage Expansion Units (#5051/5052). Additionally, disk units attached with #6502 or #6512 and not in a RAID array can be mirrored and/or unprotected. In the RAID configuration, disk unit protection is provided at less cost than mirroring, and with greater performance than system checksums. The 6502 and 6512 also supports mixing different internal disk features on the same controller.

Table C.1. 6502/6512 S	Table C.1. 6502/6512 Supported DASD Models						
DASD	Model	MB/arm	RAID	Write Cache			
6605	050	1031	No	Yes			
6605	070	1031	Yes	Yes			
6605	072	902	Yes	Yes			
6605	074	773	Yes	Yes			
6606	050	1967	No	Yes			
6606	070	1967	Yes	Yes			
6606	072	1721	Yes	Yes			
6606	074	1475	Yes	Yes			
6607	050	4194	No	Yes			
6607	070	4194	Yes	Yes			
6607	072	3669	Yes	Yes			
6607	074	3145	Yes	Yes			
6713	050	8589	No	Yes			
6713	070	8589	Yes	Yes			
6713	072	7515	Yes	Yes			
6713	074	6441	Yes	Yes			
6714	050	17548	No	Yes			
6714	070	17548	Yes	Yes			
6714	072	15354	Yes	Yes			
6714	074	13161	Yes	Yes			

6502/6512 Supported DASD Models

A minimum of four drives of the same capacity are needed to protect them with RAID-5 protection. A maximum of two arrays are allowed per controller, with a maximum of ten drives allowed per array. All drives in an array must be of the same capacity. Parity is spread on four or eight drives. 1 GB and larger disk units can be RAID-5 protected by the controller. Each System Unit Expansion Tower can support up to 16 disk units on 1 6502/6512 disk controller when the #5052 Storage Expansion Unit is installed. Each DASD Expansion Tower can support up to 32 disk units on 2 6502/6512 disk controllers.

For the Model 400, up to eight 1 GB disk units can be supported when using the #7117 Integrated Expansion Unit.

6530 Storage Device Controller

The #6530 is a Storage controller for up to 16 disk units installed in the Storage Expansion Units (#5051/5052). The 6530 does NOT have a Write Cache and does NOT support RAID. The 6530 also supports mixing different internal disk features on the same controller.

Table C.2. 6530 Supported DASD Models							
DASD	Model	MB/arm	RAID	Write Cache			
6605	030	1031	No	No			
6606	030	1967	No	No			
6607	030	4194	No	No			
6713	030	8589	No	No			
6714	030	17548	No	No			

6530 Supported DASD Models

Each System Unit Expansion Tower can support up to 16 disk units on 1 6530 disk controller when the #5052 Storage Expansion Unit is installed. Each DASD Expansion Tower can support up to 32 disk units on 2 6530 disk controllers.

6533/6754 Disk Unit Controller for RAID

Feature #6533 (also #6532) is a disk controller with a 4MB write-cache. It can provide RAID-5 protection for up to 16 internal disk units installed in the Storage Expansion Units.Feature #6754 (also #6751) is a Multi-Function IOP with a 4MB write-cache. It can provide RAID-5 protection for up to 20 internal disk units installed in the System Unit.Additionally, disk units attached with the #6533 or #6754 and not in a RAID array can be mirrored and/or unprotected.The 6533 and 6754 also support mixing different internal disk features on the same controller.

6533/6754 Supported DASD Models

Table C.3. 6533/6754 S	Table C.3. 6533/6754 Supported DASD Models							
DASD	Model	MB/arm	RAID	Write Cache				
6606	050	1967	No	Yes				
6806	050	1967	No	Yes				
6606	070	1967	Yes	Yes				
6806	070	1967	Yes	Yes				
6606	072	1721	Yes	Yes				
6806	072	1721	Yes	Yes				
6606	074	1475	Yes	Yes				
6806	074	1475	Yes	Yes				
6607	050	4194	No	Yes				
6807	050	4194	No	Yes				
6607	070	4194	Yes	Yes				
6807	070	4194	Yes	Yes				
6607	072	3669	Yes	Yes				
6807	072	3669	Yes	Yes				
6607	074	3145	Yes	Yes				
6807	074	3145	Yes	Yes				
6713	050	8589	No	Yes				
6813	050	8589	No	Yes				
6713	070	8589	Yes	Yes				
6813	070	8689	Yes	Yes				
6713	072	7515	Yes	Yes				
6813	072	7515	Yes	Yes				
6713	074	6441	Yes	Yes				
6813	074	6441	Yes	Yes				
6717	050	8589	No	Yes				
6717	070	8589	Yes	Yes				
6717	072	7515	Yes	Yes				
6717	074	6441	Yes	Yes				
6714	050	17548	No	Yes				
6714	070	14548	Yes	Yes				
6714	072	15354	Yes	Yes				
6714	074	13161	Yes	Yes				
Note: 6806, 6807, 6813	3, and 6714 are Ultra-SCS	SI DASD and 6717 is a SC	SI Wide-Ultra2 DASD.					

A minimum of four drives of the same capacity are needed to protect them with RAID-5 protection. A maximum of two arrays are allowed per controller, with a maximum of ten drives allowed per array. All drives in an array must be of the same capacity. Parity is spread on four or eight drives. Each SystemUnit Expansion Tower can support up to 16 disk units on 1 6533 disk controller. Each DASD Expansion Tower can support up to 32 disk units on 2 6533 disk controllers.

For the Models 640 and 650, up to 20 disk units can be supported on 1 6754 MFIOP.

2741/2740/2726 Disk Unit Controller for RAID

Feature #2741 (also #2726) is a disk controller with a 4MB write-cache. It can provide RAID-5 protection for up to 15 internal disk units installed in the PCI System Unit, PCI Expansion Unit or PCI Expansion Tower. Feature #2740 is a low cost disk controller with a 4MB write-cache that is targeted for smaller systems. It can provide RAID-5 protection for up to 10 internal disk units installed in the PCI System Unit or PCI Expansion Unit. Additionally, disk units attached with the #2741 or #2740 and not in a RAID array

can be mirrored and/or unprotected. The 2741 and 2740 also support mixing different internal disk features on the same controller.

DASD	Model	MB/arm	RAID	Write Cache
6606	050	1967	No	Yes
6806	050	1967	No	Yes
6606	070	1967	Yes	Yes
6806	070	1967	Yes	Yes
6606	072	1721	Yes	Yes
6806	072	1721	Yes	Yes
6606	074	1475	Yes	Yes
6806	074	1475	Yes	Yes
6607	050	4194	No	Yes
6807	050	4194	No	Yes
6607	070	4194	Yes	Yes
6807	070	4194	Yes	Yes
6607	072	3669	Yes	Yes
6807	072	3669	Yes	Yes
6607	074	3145	Yes	Yes
6807	074	3145	Yes	Yes
6713	050	8589	No	Yes
6813	050	8589	No	Yes
6713	070	8689	Yes	Yes
6813	070	8589	Yes	Yes
6713	072	7515	Yes	Yes
6813	072	7515	Yes	Yes
6713	074	6441	Yes	Yes
6813	074	6441	Yes	Yes
6717	050	8589	No	Yes
6717	070	8589	Yes	Yes
6717	072	7515	Yes	Yes
6717	074	6441	Yes	Yes
6714	050	17548	No	Yes
6714	070	17548	Yes	Yes
6714	072	15354	Yes	Yes
6714	074	13161	Yes	Yes

2741/2740/2726 Supported DASD Models

A minimum of four drives of the same capacity are needed to protect them with RAID-5 protection. A maximum of two arrays are allowed per controller, with a maximum of ten drives allowed per array. All drives in an array must be of the same capacity. Parity is spread on four or eight drives.

2748 PCI RAID Unit Controller

Feature #2748 is a PCI disk controller with a 26MB write-cache. It can provide RAID-5 protection for up to 15 internal disk units installed in the PCI System Unit, PCI Expansion Unit, or PCI Expansion Tower. Additionally, disk units attached with the #2748 and not in a RAID array can be mirrored and/or unprotected. The 2748 also supports mixing different internal disk features on the same controller. When DASD Compresson is enabled, the write-cache is limited to 4MB.

2748 Supported DASD Models

Table C.5. 2748 Suppo	Table C.5. 2748 Supported DASD Models						
DASD	Model	MB/arm	RAID	Write Cache			
6606	050	1967	No	Yes			
6806	050	1967	No	Yes			
6606	070	1967	Yes	Yes			
6806	070	1967	Yes	Yes			
6606	072	1721	Yes	Yes			
6806	072	1721	Yes	Yes			
6606	074	1475	Yes	Yes			
6806	074	1475	Yes	Yes			
6607	050	4194	No	Yes			
6807	050	4194	No	Yes			
6607	070	4194	Yes	Yes			
6807	070	4194	Yes	Yes			
6607	072	3669	Yes	Yes			
6807	072	3669	Yes	Yes			
6607	074	3145	Yes	Yes			
6807	074	3145	Yes	Yes			
6713	050	8589	No	Yes			
6813	050	8589	No	Yes			
6713	070	8689	Yes	Yes			
6813	070	8589	Yes	Yes			
6713	072	7515	Yes	Yes			
6813	072	7515	Yes	Yes			
6713	074	6441	Yes	Yes			
6813	074	6441	Yes	Yes			
6717	050	8589	No	Yes			
6717	070	8589	Yes	Yes			
6717	072	7515	Yes	Yes			
6717	074	6441	Yes	Yes			
6714	050	17548	No	Yes			
6714	070	17548	Yes	Yes			
6714	072	15354	Yes	Yes			
6714	074	13161	Yes	Yes			
Note: 6806, 6807, 6813	3, and 6714 are Ultra-SC	SI DASD and 6717 is a SC	SI Wide-Ultra2 DASD.				

A minimum of four drives of the same capacity are needed to protect them with RAID-5 protection. A maximum of two arrays are allowed per controller, with a maximum of ten drives allowed per array. All drives in an array must be of the same capacity. Parity is spread on four or eight drives.

9728 Storage Device Controller

The #9728 is a Storage controller for up to 5 disk units installed in the PCI System Unit. The 9728 does NOT have a Write Cache and does NOT support RAID. The 9728 also supports mixing different internal disk features on the same controller.

9728 Supported DASD Models

Table C.6. 9728 Suppo				
DASD	Model	MB/arm	RAID	Write Cache
6606	030	1967	No	Yes
6806	030	1967	No	Yes
6607	030	4194	No	Yes
6807	030	4194	No	Yes
6713	030	8589	No	Yes
6813	030	8589	No	Yes
6714	030	17548	No	No
Note: 6806, 6807, 6813	3, and 6714 are Ultra-SCS	I DASD		

Appendix D. AS/400 CPW Values

This chapter details the system capacities based on a workload called Commercial Processing Workload (CPW). For a detailed description, refer to *Appendix A*, "*CPW Benchmark Description*".

CPW values are relative system performance metrics and reflect the relative system capacity for the CPW workload. CPW values can be used with caution in a capacity planning analysis (e.g., to scale CPU-constrained capacities, CPU time per transaction). However, these values may not properly reflect specific workloads other than CPW because of differing detailed characteristics (e.g., cache miss ratios, average cycles per instruction, software contention).

Additional factors that determine what performance is achievable are: type and number of disk devices, number of work station controllers, amount of memory, system model, and the application being run.

The following tables show for each server model the maximum interactive CPW and its corresponding CPU % and the point (the knee of the curve) where the interactive utilization begins to increasingly impact client/server performance. For the models that have multiple processors, and the knee of the curve is also given in CPU%, the percent value is the percent of all the processors (not of a single one). CPW values may be increased as enhancements are made to the operating system (e.g. each feature of the Model 53S for V3R7 and V4R1). The server model behavior is fixed to the original CPW values. For example, the model 53S-2157 had V3R7 CPWs of 509.9/30.7 and V4R1 CPWs 650.0/32.2. When using the 53S with V4R1, this means the knee of the curve is 2.6% CPU and the maximum interactive is 7.7% CPU, the same as it was in V3R7.

The CPW values shown in the tables are based on IBM internal tests. Actual performance in a customer environment will vary.

Table values in bold indicate published CPW values.

For additional CPW values, see the IBM AS/400 Advanced 36 Performance Capabilities Reference.

D.1 V4R4 Additions

The Model 7xx is new in V4R4. Also in V4R4 are the Model 170s features 2289 and 2388 were added. See the chapter, **AS/400 RISC Server Model Performance Behavior,** for a description of the performance highlights of these new models.

Testing in the Rochester laboratory has shown that for systems executing traditional commercial applications such as RPG or COBOL interactive general business applications may experience about a 5% increase in CPU requirements. This effect was observed using the workload used to compute CPW, as shown in the tables that follows. Except for systems which are nearing the need for an upgrade, we do not expect this increase to significantly affect transaction response times. It is recommended that other sections of the Performance Capabilities Reference Manual (or other sizing and positioning documents) be used to estimate the impact of upgrading to the new release.

D.2 AS/400e Model 7xx Servers

MAX Interactive CPW = Interactive CPW (Knee) * 7/6 CPU % used by Interactive @ Knee = Interactive CPW (Knee) / Processor CPW * 100 CPU % used by Processor @ Knee = 100 - CPU % used by Interactive @ Knee CPU % used by Interactive @ Max = Max Interactive CPW / Processor CPW * 100

Table D.2.1 Model /XX S	Servers (a	all new Norths	tar model	ls)		
Model C	Thip Speed	L2 cache	CPUs	Processor CPW	Interactive CPW	Interactive CPW
	MHz	per CPU			(Knee)	(Max)
720-2061 (Base)	200	n/a	1	240	35	40.8
720-2061 (1501)	200	n/a	1	240	70	81.7
720-2061 (1502)	200	n/a	1	240	120	140
720-2062 (Base)	200	4 MB	1	420	35	40.8
720-2062 (1501)	200	4 MB	1	420	70	81.7
720-2062 (1502)	200	4 MB	1	420	120	140
720-2062 (1503)	200	4 MB	1	420	240	280
720-2063 (Base)	200	4 MB	2	810	35	40.8
720-2063 (1502)	200	4 MB	2	810	120	140
720-2063 (1503)	200	4 MB	2	810	240	280
720-2063 (1504)	200	4 MB	2	810	560	653.3
720-2064 (Base)	255	4 MB	4	1600	35	40.8
720-2064 (1502)	255	4 MB	4	1600	120	140
720-2064 (1503)	255	4 MB	4	1600	240	280
720-2064 (1504)	255	4 MB	4	1600	560	653.3
720-2064 (1505)	255	4 MB	4	1600	1050	1225
730-2065 (Base)	262	4 MB	1	560	70	81.7
730-2065 (1507)	262	4 MB	1	560	120	140
730-2065 (1508)	262	4 MB	1	560	240	280
730-2065 (1509)	262	4 MB	1	560	560	653.3
720,2066 (Pasa)	262	4 MP	2	1050	70	<u> </u>
730-2066 (1507)	202	4 MB	2	1050	120	140
730 2066 (1507)	262	4 MB	2	1050	240	280
730-2000 (1508)	202	4 MB	2	1050	<u> </u>	653.3
730-2066 (1510)	262	4 MB	2	1050	1050	1225
		() ()				01.5
730-2067 (Base)	262	4 MB	4	2000	70	81.7
730-2067 (1508)	262	4 MB	4	2000	240	280
730-2067 (1509)	262	4 MB	4	2000	560	653.3
730-2067 (1510)	262	4 MB	4	2000	1050	1225
730-2067 (1511)	262	4 MB	4	2000	2000	2333.3
730-2068 (Base)	262	4 MB	8	2890	70	81.7
730-2068 (1508)	262	4 MB	8	2890	240	280
730-2068 (1509)	262	4 MB	8	2890	560	653.3
730-2068 (1510)	262	4 MB	8	2890	1050	1225
730-2068 (1511)	262	4 MB	8	2890	2000	2333.3
740-2069 (Base)	262	8 MB	8	3660	120	140
740-2069 (1510)	262	8 MB	8	3660	1050	1225
740-2069 (1511)	262	8 MB	8	3660	2000	2333.3
740-2069 (1512)	262	8 MB	8	3660	3660	4270
740-2070 (Base)	262	8 MB	12	4550	120	140
740-2070 (1510)	262	8 MR	12	4550	1050	1225
740-2070 (1511)	262	8 MB	12	4550	2000	2333.3
740-2070 (1512)	262	8 MB	12	4550	3660	4270
740-2070 (1513)	262	8 MB	12	4550	4550	5308.3

D.3 Model 170 Servers

Current 170 Servers

MAX Interactive CPW = Interactive CPW (Knee) * 7/6 CPU % used by Interactive @ Knee = Interactive CPW (Knee) / Processor CPW * 100 CPU % used by Processor @ Knee = 100 - CPU % used by Interactive @ Knee CPU % used by Interactive @ Max = Max Interactive CPW / Processor CPW * 100

Table D.3.1 Current Model 170 Servers (Note: 2289 and 2388 are two ne)					two new North	star models)			
Feature #	CPUs	Chip Speed	L2 cache per CPU	Processor CPW	Interactive CPW (Knee)	Interactive CPW (Max)	Processor CPU % @ Knee	Interactive CPU % @ Knee	Interactive CPU % @ Max
2289	1	200 MHz	n/a	50	15	17.5	70	30	35
2290	1	200 MHz	n/a	73	20	23.3	72.6	27.4	32
2291	1	200 MHz	n/a	115	25	29.2	78.3	21.7	25.4
2292	1	200 MHz	n/a	220	30	35	86.4	13.6	15.9
2385	1	252 MHz	4 MB	460	50	58.3	89.1	10.9	12.7
2386	1	252 MHz	4 MB	460	70	81.7	84.8	15.2	17.8
2388	2	255 MHz	4 MB	1090	70	81.7	92.3	6.4	7.5

Note: the CPU not used by the interactive workloads at their Max CPW is used by the system CFINTnn jobs. For example, for the 2386 model the interactive workloads use 17.8% of the CPU at their maximum and the CFINTnn jobs use the remaining 82.2%. The processor workloads use 0% CPU when the interactive workloads are using their maximum value.

AS/400e Dedicated Server for Domino

Table D.3	3.2 Dedic	ated Server	for Domine	0					
		Chin	I 2 cacha	Processor		Processor	Processor	Interactive	Interactive
Feature #	CPUs	Speed	D2 Cache	CDW	Interactive	CPU%	CPU %	CPU %	CPU %
		speed	per CFU	Crw	CPW	@ Knee	@ Max	@ Knee	@ Max
2407	1	n/a	n/a	30	10	-	-	-	-
2408	1	n/a	4 MB	60	15	-	-	-	-
2409	2	n/a	4 MB	120	20	-	-	-	_

Previous Model 170 Servers

On previous Model 170's the knee of the curve is about 1/3 the maximum interactive CPW value.

Note that a constrained (c) CPW rating means the maximum memory or DASD configuration is the constraining factor, not the processor. An unconstrained (u) CPW rating means the processor is the first constrained resource.

Table D.3.2	Table D.3.2 Previous Model 170 Servers										
Footuro #	Constrain /	Client / Server	Interactive	Interactive CPW	Interactive	Interactive					
reature #	Unconstr	CPW	CPW (Max)	(Knee)	CPU % @ Max	CPU % @ Knee					
2159	с	73	16	5.3	22.2	7.7					
	u	73	16	5.3	22.2	7.7					
2160	с	114	23	7.7	21.2	7.4					
	u	114	23	7.7	21.2	7.4					

Feature #	Constrain / Unconstr	Client / Server CPW	Interactive CPW (Max)	Interactive CPW (Knee)	Interactive CPU % @ Max	Interactive CPU % @ Knee
2164	с	125	29	9.7	14	4.7
2164	u	210	29	9.7	14	4.7
2176	с	125	40	13.3	12.9	4.4
2170	u	319	40	13.3	12.9	4.4
2183	с	125	67	22.3	21.5	7.2
	u	319	67	22.3	21.5	7.2

D.4 AS/400e Servers

Table D.4.1 AS/400e Servers										
Model	Easture #	CDU	Max	Max	1/3 Max	CPU % @ Max	CPU %			
Model	reature #	CPUS	C/S CPW	Inter CPW	Interact CPW	Interact	@ the Knee			
\$10	2118	1	45.4	16.2	5.4	35.7	11.9			
510	2119	1	73.1	24.4	8.1	33.4	11.1			
	2161	1	113.8	31	10.3	27.2	9.1			
S20	2163	1	210	35.8	11.9	17	5.7			
	2165	2	464.3	49.7	16.7	10.7	3.6			
	2166	4	759	56.9	19.0	7.5	2.5			
	2257	1	319	51.5	17.2	16.1	5.4			
\$20	2258	2	583.3	64	21.3	11	3.7			
530	2259	4	998.6	64	21.3	6.4	2.1			
	2260	8	1794	64	21.3	3.6	1.2			
	2207	8	3660	120	40	3.2	1.1			
540	2208	12	4550	120	40	2.6	0.8			
540	2256	8	1794	64	21.3	3.6	1.2			
	2261	12	2340	64	21.3	2.7	0.9			

For AS/400e servers the knee of the curve is about 1/3 the maximum interactive CPW value.

D.5 AS/400e Custom Servers

For custom servers the knee of the curve is about 6/7 maximum interactive CPW value.

Table D.5.1 AS/400e Custom Servers									
Model	Feature #	CPUs	Max	Max	6/7 Max	CPU % @	CPU %		
520	2177	4	759	110.7	94.9	14.6	12.5		
520	2178	4	759	221.4	189.8	29.2	25.0		
	2320	4	998.6	215.1	184.4	21.5	18.5		
S30	2321	8	1794	386.4	331.2	21.5	18.5		
	2322	8	1794	579.6	496.8	32.5	27.7		
S40	2340	8	3660	1050.0	900.0	28.6	24.5		
	2341	12	4550	2050.0	1757.1	38.6	33.1		

D.6 AS/400 Advanced Servers

For AS/400 Advanced Servers the knee of the curve is about 1/3 the maximum interactive CPW value.

For releases prior to V4R1 the model 150 was constrained due to the memory capacity. With the larger capacity for V4R1, memory is no longer the limiting resource. In V4R1, the limit of 4 DASD devices is the constraining resource. For workloads that do not perform as many disk operations or don't require as much memory, the unconstrained CPW value may be more representative of the performance capabilities. An unconstrained CPW rating means the processor is the first constrained resource.

Table D	.6.1 AS/400	Advanced Ser	vers: V4	R1 and V4R2				
Model	Eastura #	Constrain /	CDU	Max	Max	1/3 Max	CPU % @ Max	CPU % @ the
Widder	Feature #	Unconstr	CrUs	C/S CPW	Inter CPW	Interact CPW	Interact	Knee
	2269	с	1	20.2	13.8	4.6	51.1	17
150	2269	u	1	27	13.8	4.6	51.1	17
150	2270	с	1	20.2	20.2	6.7	61.9	20.6
	2270	u	1	35	20.6	6.9	61.9	20.6
405	2109	n/a	1	27	9.4	3.1	30.1	10
405	2110	n/a	1	35	14.5	3.9	37.4	12.5
	2111	n/a	1	63.0	21.6	7.2	29.8	9.9
	2112	n/a	1	91.0	32.2	10.8	29.8	9.9
50S	2120	n/a	1	81.6	22.5	8.1	27.8	9.3
	2121	n/a	1	111.5	32.2	10.7	30	10
	2122	n/a	1	138.0	32.2	12.0	23.8	8.9
	2154	n/a	1	188.2	32.2	15.9	20.3	6.8
53S	2155	n/a	2	319.0	32.2	10.7	13.5	4.5
	2156	n/a	4	598.0	32.2	10.7	9	3
	2157	n/a	4	650.0	32.2	10.9	7.7	2.6

Table D	Table D.6.2 AS/400 Advanced Servers: V3R7										
Model	Footuro #	Constrain /	CPU	Max	Max	1/3 Max	CPU % @ Max	CPU %			
Widder	Teature #	Unconstr	CIUS	C/S CPW	Inter CPW	Interact CPW	Interact	@ the Knee			
	2269	с	1	10.9	10.9	3.6	100.0	33.0			
150	2269	u	1	10.9	10.9	3.6	100.0	33.0			
150	2270	с	1	27.0	13.8	4.6	51.1	17.0			
	2270	u	1	33.3	20.6	6.9	61.9	20.6			
	2109	n/a	1	27.0	9.4	3.1	30.1	10			
40S	2110	n/a	1	33.3	13.8	3.7	37.4	12.5			
	2111	n/a	1	59.8	20.6	6.9	29.8	9.9			
	2112	n/a	1	87.3	30.7	10.3	29.8	9.9			
505	2120	n/a	1	77.7	21.4	7.7	27.8	9.3			
505	2121	n/a	1	104.2	30.7	10.2	30	10			
	2122	n/a	1	130.7	30.7	11.5	23.8	8.9			
	2154	n/a	1	162.7	30.7	13.3	20.3	6.8			
528	2155	n/a	2	278.8	30.7	10.2	13.5	4.5			
555	2156	n/a	4	459.3	30.7	10.2	9	3			
	2157	n/a	4	509.9	30.7	10.4	7.7	2.6			

D.7 AS/400e Custom Application Server Model SB1

AS/400e application servers are particularly suited for environments with minimal database needs, minimal disk storage needs, lots of low-cost memory, high-speed connectivity to a database server, and minimal upgrade importance.

The throughput rates for Financial (FI) dialogsteps (ds) per hour may be used to size systems for customer orders. Note: 1 SD ds = = 2.5 Fl ds. (SD = Sales & Distribution).

Table D.7.1 AS/400e Custom Application Server Model SB1									
Model	CPUs	SAP SD ds/hr		FI ds/hr					
	01 05	Release	@ 65% CPU Utilization	@ 65% CPU Utilization					
2212	0	3.1H	109,770.49	274,426.23					
2512	8	4.0B	65,862.29	164,655.74					
2212	12	3.1H	158,715.76	396,789.40					
2313		4.0B	95,229.46	238,073.64					

D.8 Previous AS/400 RISC System Capacities

Table D.8	.1 AS/400 RISC S	ystems				
Model	Feature Code	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	V3R7 CPW	V4R1 CPW
	2130	1	160	50	13.8	13.8
400	2131	1	224	50	20.6	20.6
400	2132	1	224	50	27	27
	2133	1	224	50	33.3	35
	2140	1	768	652	21.4	21.4
500	2141	1	768	652	30.7	30.7
	2142	1	1024	652	43.9	43.9
510	2143	1	1024	652	77.7	81.6
510	2144	1	1024	652	104.2	111.5
	2150	1	4096	996	131.1	148
	2151	1	4096	996	162.7	188.2
530	2152	2	4096	996	278.8	319
	2153	2	4096	996	459.3	598
	2162	4	4096	996	509.9	650

Table D.8	.2 AS/400e System	S			
Model	Feature Code	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	V4R3 CPW
	2129	1	384	175.4	22.7
600	2134	1	384	175.4	32.5
000	2135	1	384	175.4	45.4
	2136	1	512	175.4	73.1
	2175	1	1856	944.8	50
	2179	1	2048	944.8	85.6
620	2180	1	2048	944.8	113.8
	2181	1	2048	944.8	210
	2182	2	4096	944.8	464.3
	2237	1	16384	1340	319
640	2238	2	8704	1340	583.3
	2239	4	16384	1340	998.6
	2188	8	40960	2095.9	3660
650	2189	12	40960	2095.9	4550
050	2240	8	32768	2095.9	1794
	2243	12	32768	2095.9	2340

D.9 AS/400 CISC Model Capacities

Table D.9.1 A	Table D.9.1 AS/400 CISC Model: 9401									
Model	Feature	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	CPW					
P02	n/a	1	16	2.1	7.3					
	2114	1	24	2.99	7.3					
P03	2115	1	40	3.93	9.6					
	2117	1	56	3.93	16.8					

Table D.9.2 AS/4	Table D.9.2 AS/400 CISC Model: 9402 Systems										
Model	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	CPW							
C04	1	12	1.3	3.1							
C06	1	16	1.3	3.6							
D02	1	16	1.2	3.8							
D04	1	16	1.6	4.4							
E02	1	24	2.0	4.5							
D06	1	20	1.6	5.5							
E04	1	24	4.0	5.5							
F02	1	24	2.1	5.5							
F04	1	24	4.1	7.3							
E06	1	40	7.9	7.3							
F06	1	40	8.2	9.6							

Table D.9.3 AS/400 CISC Model: 9402 Servers							
Feature Code	CPUs	PUs Memory (MB) Maximum Disk (GB) Maximum C/S CPW In					
S01	1	56	3.9	17.1	5.5		
100	1	56	7.9	17.1	5.5		

Table D.9.4 AS/400 CISC Model: 9404 Systems							
Model	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	CPW			
B10	1	16	1.9	2.9			
C10	1	20	1.9	3.9			
B20	1	28	3.8	5.1			
C20	1	32	3.8	5.3			
D10	1	32	4.8	5.3			
C25	1	40	3.8	6.1			
D20	1	40	4.8	6.8			
E10	1	40	19.7	7.6			
D25	1	64	6.4	9.7			
F10	1	72	20.6	9.6			
E20	1	72	19.7	9.7			
F20	1	80	20.6	11.6			
E25	1	80	19.7	11.8			
F25	1	80	20.6	13.7			

Table D.9.5 AS/400 CISC Model: 9404 Servers							
Feature Code	CPUs	Memory (MB) Maximum Disk (GB) Maximum C/S CPW Inte					
135	1	384	27.5	32.3	9.6		
140	2	512	47.2	65.6	11.6		

Table D.9.6 AS/400 CISC Model: 9406 Systems						
Model	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	CPW		
B30	1	36	13.7	3.8		
B35	1	40	13.7	4.6		
B40	1	40	13.7	5.2		
B45	1	40	13.7	6.5		
D35	1	72	67.0	7.4		
B50	1	48	27.4	9.3		
E35	1	72	67.0	9.7		
D45	1	80	67.0	10.8		
D50	1	128	98.0	13.3		
E45	1	80	67.0	13.8		
F35	1	80	67.0	13.7		
B60	1	96	54.8	15.1		
F45	1	80	67.0	17.1		
E50	1	128	98.0	18.1		
B70	1	192	54.8	20.0		
D60	1	192	146	23.9		
F50	1	192	114	27.8		
E60	1	192	146	28.1		
D70	1	256	146	32.3		
E70	1	256	146	39.2		
F60	1	384	146	40.0		
D80	2	384	256	56.6		
F70	1	512	256	57.0		
E80	2	512	256	69.4		
E90	3	1024	256	96.7		
F80	2	768	256	97.1		
E95	4	1152	256	116.6		
F90	3	1024	256	127.7		
F95	4	1280	256	148.8		
F97	4	1536	256	177.4		

Table D.9.7 AS/400 Advanced Systems (CISC)						
Model	Feature Code	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	CPW	
	2030	1	24	23.6	7.3	
200	2031	1	56	23.6	11.6	
	2032	1	128	23.6	16.8	
	2040	1	72	117.4	11.6	
300	2041	1	80	117.4	16.8	
	2042	1	160	117.4	21.1	
310	2043	1	832	159.3	33.8	
	2044	2	832	159.3	56.5	
320	2050	1	1536	259.6	67.5	
	2051	2	1536	259.6	120.3	
	2052	4	1536	259.6	177.4	

Table D.9.8 AS/400 Advanced Servers (CISC)							
Model	Feature Code	CPUs	Memory (MB) Maximum	Disk (GB) Maximum	C/S CPW	Interactive CPW	
20S	2010	1	128	23.6	17.1	5.5	
2FS	2010	1	128	7.8	17.1	5.5	
2SG	2010	1	128	7.8	17.1	5.5	
2SS	2010	1	128	7.8	17.1	5.5	
305	2411	1	384	86.5	32.3	9.6	
	2412	2	832	86.5	68.5	11.6	