




The System Administrator's Companion to AS/400 Availability and Recovery

*Susan Powers Ellen Dreyer Andersen
Rob Jones Hubert Lye Petri Nuutinen*



International Technical Support Organization

<http://www.redbooks.ibm.com>

This book was printed at 240 dpi (dots per inch). The final production redbook with the RED cover will be printed at 1200 dpi and will provide superior graphics resolution. Please see "How to Get ITSO Redbooks" at the back of this book for ordering instructions.



International Technical Support Organization

SG24-2161-00

**The System Administrator's Companion
to AS/400 Availability and Recovery**

August 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix F, "Special Notices" on page 411.

First Edition (August 1998)

This edition applies to Version 4 Release 2 Modification 0 and prior of IBM OS/400 Operating System 5769-SS1

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1998. 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

Preface	xiii
The Team That Wrote This Redbook	xiii
Comments Welcome	xvi
Chapter 1. Introduction to AS/400 Availability and Recovery	1
1.1 A Historical Perspective of AS/400 Availability Enhancements	2
1.2 Frequently Asked Questions about Availability and Recovery	6
1.2.1 Disk Management	6
1.2.2 Database Management	6
1.2.3 Saving or Restoring the System	7
1.2.4 System Availability Factors	9
1.2.5 System Management	10
1.2.6 Communications or Network Issues	12
Chapter 2. Availability and Recovery Concepts	13
2.1 The Importance of Backup	14
2.2 Backup and Recovery Planning	14
2.3 Levels of Availability	17
2.4 Types of Unplanned Outages	18
2.4.1 Power Failure	19
2.4.2 Unprotected DASD or Disk Failure	19
2.4.3 System Failure	19
2.4.4 Human Error or Program Failure	19
2.4.5 Site Loss	19
2.5 Recovery Steps	20
Chapter 3. Availability Options Provided by Hardware	23
3.1 Load Source Protection	23
3.1.1 Mirrored Protection	24
3.1.2 Standard Mirrored Protection	24
3.1.3 Mirrored Load Source Protection Prior to V3R7	25
3.1.4 Remote Load Source Mirrored Protection on V3R7 and Later Systems	25
3.2 Device Parity Protection	27
3.3 Uninterruptible Power Supply	28
3.4 Battery Backup	29
3.5 Continuously Powered Main Storage	29
3.6 Tape Device Options	30
3.7 Alternate Installation Device	30
Chapter 4. IPL Improvements for Availability	33
4.1 A Basic Understanding of an IPL	33
4.2 IPL Types	34
4.3 Affecting the Time to IPL	36
4.4 System Managed Access Path Protection	38
4.4.1 SMAPP Tasks	39
4.4.2 Performance Considerations when SMAPP is Activated	40
4.4.3 Modifying SMAPP	40
4.5 Changing IPL Attributes	42
4.5.1 Restart Type (RESTART) Parameter	43
4.5.2 Hardware Diagnostics (HDWDIAG) Parameter	43

4.5.3 Check Job Tables	44
4.6 Marking the Progress of an IPL with SRC Codes	44
4.7 IPL Benchmarks	46
4.7.1 Benchmark Configuration	47
Chapter 5. Save and Restore for Availability and Recovery	49
5.1 SAVxxx and RSTxxx and Flexibility	49
5.1.1 Concurrent Save with Generic OMITLIB Example	50
5.2 Save While Active and Object Locks	50
5.2.1 Save While Active Considerations	53
5.2.2 Save While Active and Target Release	53
5.3 Omitting Objects on a SAVSYS Operation	53
5.4 Concurrent Save Operations	55
5.4.1 Concurrent Saves on Libraries	55
5.4.2 Concurrent Saves for DLOs	56
5.5 Concurrent Restore Operations	57
5.6 Use Optimum Blocking for Save and Restore	57
5.7 Save and Restore Scenario	58
5.7.1 Save Strategy	59
5.7.2 Additional Considerations	61
5.8 Unattended Saves Using the SAVE Menu	61
5.8.1 Start Time	63
5.8.2 V4R2 Availability Options on the SAVE Menu	63
5.9 ObjectConnect/400	65
5.9.1 ObjectConnect/400 Command Sets	66
5.9.2 ObjectConnect/400 Offers Simplicity	67
5.9.3 ObjectConnect/400 Implementation Considerations	69
5.10 Save and Restore Spooled Files	69
5.10.1 QUSRTOOL to Save Spooled Files	69
5.10.2 Creating CL Commands to Save Spooled Files	70
5.11 Multinational Environments and Object Names	72
5.12 Product Preview for Save Restore Enhancements	73
Chapter 6. Save and Restore Considerations for Mixed Release Environments	75
6.1 Target Release and Save While Active	75
6.2 Observability Considerations when Restoring Objects	77
6.3 USEOPTBLK for Save Performance	78
6.3.1 USEOPTBLK and TGTRLS Compatibility	79
6.3.2 USEOPTBLK and DTACPR Compatibility	80
6.3.3 USEOPTBLK and Other Considerations	80
6.3.4 Correcting a Back Level QUSRSYS and QGPL Library on a CISC to RISC Migration	81
6.4 Journal Receivers and Previous Release Systems	83
Chapter 7. Licensed Program and PRPQ Backup and Recovery	85
7.1 Licensed Program and PRPQ Considerations	85
7.2 Saving Licensed Programs	86
7.3 Restoring Licensed Programs	87
7.4 Special Considerations to Install and Restore Licensed Programs	88
7.4.1 LICPGM Menu Does Not Manage all Licensed Programs	88
7.4.2 User Profile Authority with Licensed Programs	89
7.4.3 Multi-National Considerations with Directory Names	89
7.5 Licensed Program Library Names for IBM Supplied Libraries	89
7.6 Restoring Commands from Licensed Program Libraries to QSYS	92
7.7 For More Information	93

Chapter 8. Save, Restore, and System Performance for Availability	95
8.1 Save and Restore Performance	95
8.1.1 Data Compaction for Tape	95
8.1.2 Use Optimum Block Size	97
8.1.3 Data Compression	97
8.1.4 Save and Restore Tips and Techniques for Better Performance	98
8.2 Defective Device or Media Considerations	100
8.3 What a User Can Do to Influence System Performance	100
8.3.1 Automatic Performance Adjustments—It Is Worth Another Look	100
8.3.2 Altering Shared Pool and Priority Values with the Automatic Tuner	101
8.3.3 Tuning the System Tuner	104
8.3.4 Tuning Parameters for Batch	104
8.3.5 Dynamic Priority and Controlling CPU Intensive Jobs	105
8.4 For More Information	107
Chapter 9. Tools for Automating System Management Functions	109
9.1 ADSTAR Distributed Storage Manager/400 (ADSM/400)	109
9.1.1 ADSM/400 Customer Scenario	110
9.1.2 ADSM Elements	110
9.1.3 Administrative Client Functions	112
9.1.4 Server Functions	112
9.1.5 Total Disaster Recovery with ADSM/400 Version 2	112
9.1.6 More ADSM/400 Information	113
9.2 ADSM/400 and BRMS/400 Interoperability	114
9.3 Backup, Recovery, and Media Services for AS/400 (BRMS)	114
9.3.1 BRMS Recovery Report	115
9.3.2 BRMS User Exits and Message Handling	115
9.3.3 BRMS Maintenance Activities—STRMNTBRM	116
9.3.4 BRMS BRMLOG	117
9.3.5 Media Classes—CPYMEDIBRM	117
9.3.6 Network Time Zone Synchronization	117
9.3.7 Large Tape File Sequence Numbers for Non-Save and Restore Operations	119
9.3.8 Shorten the Time to Save with USEOPTBLK	119
9.3.9 SAVSYSBRM Command Support for OMIT Parameter	119
9.4 WRKASP Utility	120
9.4.1 ASP Monitoring	121
9.4.2 ASP Configuration	121
9.4.3 ASP Management	122
9.5 Print System Information Tool	123
9.6 OS/400 Job Scheduler (Part of the Operating System)	125
9.6.1 Work with Job Schedule Entries	125
9.7 IBM Job Scheduler for OS/400	127
9.7.1 Job Scheduling and Availability	127
9.8 IBM SystemView System Manager for AS/400	128
9.9 IBM SystemView Managed System Services for AS/400	130
9.9.1 Scenario of Using SM/400 with MSS/400	131
9.10 Automating Message Management	131
9.10.1 QSYSMSG Message Queue	131
9.10.2 Break Handling Program (User-Exit Program)	132
9.10.3 System Reply List	134
9.10.4 Operational Assistant	134
9.11 OS/400 Alert Support	135
9.12 Automating Security Management	137
9.12.1 Audit Journal	138

Chapter 10. Work Management for System Availability	139
10.1 Work With Active Jobs	139
10.2 Work Control Block Table	141
10.2.1 Work Control Block Table Cleanup	143
10.3 Display Job Tables	143
10.3.1 Permanent Job Structures Field	144
10.3.2 Temporary Job Structures Field	145
10.3.3 Entries Field	145
10.4 System Jobs Affecting Availability	146
10.4.1 QSYSARB and QSYSARBn Jobs	147
10.4.2 QJOBSCD—Job Scheduler	148
10.5 System Values Affecting Availability	148
10.5.1 Auxiliary Storage Lower Limit—QSTGLOWLMT	149
10.5.2 Auxiliary Storage Lower Limit Action	150
10.5.3 Device Recovery Action	150
10.6 QSYSOPR Message Queue Wrap When Full	151
10.7 When CPM or Dump Processing Hang	152
10.8 When the System Date is Reset	152
10.9 End Job Abnormal	152
10.10 Reclaim Storage	153
10.10.1 The Benefits of Running RCLSTG	154
10.10.2 Reclaim Storage Options	155
10.10.3 Reclaim Storage Status Messages	156
10.10.4 Reclaim Storage Error Messages	157
10.10.5 Completion Time for Reclaim Storage	157
10.11 Other Reclaim Processes	158
10.11.1 Reclaim Document Library Object (RCLDLO)	158
10.11.2 Reclaim Spool Storage (RCLSPLSTG)	159
10.11.3 Detecting Damage in QSYS Physical Files	159
10.11.4 Detecting Damage in Physical Files	160
10.12 Power Down System	161
10.12.1 Change Command Default Considerations	161
10.12.2 End Subsystem Option	162
10.12.3 Timeout Options for Power Down System	162
10.12.4 Time to Terminate Improves System Availability	164
10.13 Work Management APIs	164
10.13.1 QUSRJOBI Retrieve Job Information API	164
10.13.2 QUSLJOB List Job API	165
10.13.3 QWTCHGJB Change Job API	165
10.13.4 QUSCHGPA Change Pool Attributes API	166
Chapter 11. Availability and the PTF Process	167
11.1 PTF Terminology	167
11.2 Preventive Service Planning	169
11.3 Applying and Removing PTFs and IPLs	170
11.3.1 LIC PTF Apply and System Availability	171
11.3.2 Applying PTFs without an IPL	171
11.4 Product Level Support	171
11.4.1 DB2/400 PTF Strategy	175
11.5 Conditional PTFs	175
11.6 Cumulative PTF Package	176
11.6.1 Applying CUM Packages	177
11.7 Applying PTFs for the Next Release Prior to Installing the Next Release	177
11.8 Distributing PTFs	177
11.9 Requesting PTFs and Cover Letters Using the Internet	178

11.9.1 Media or PTF Cover Letter Order Scenario	179
11.10 For More Information	179
Chapter 12. Communications Error Recovery and Availability	181
12.1 What Communications Error Recovery Procedures Are	181
12.2 Improvements on V4R2 Systems	181
12.2.1 QCMNARB System Value Setting	182
12.2.2 High-Performance Routing	183
12.2.3 Device Recovery Performance for Display Devices	185
12.2.4 MAXFRAME Value on LAN Controller Descriptions	185
12.2.5 Force a Vary Off	186
12.2.6 LAN Response Timer	186
12.2.7 Device Allocation Messages	187
12.2.8 QAUTOVRT System Value	188
12.2.9 Removal of Obsolete Messages in QSYSOPR	188
12.2.10 TCP Error Recovery	188
12.2.11 Serviceability Improvements	190
12.3 Improvements on V4R1 Systems	191
12.3.1 Restructure of 5250 Display Station Pass Through	191
12.3.2 File Server Job Restructure	197
12.3.3 Activation of the Operational LAN Manager	198
12.3.4 Program Start Request Message Threshold	198
12.3.5 Error Log Filtering	199
12.3.6 Communications Trace Improvements	199
12.4 Configuration Tips and Techniques	200
12.4.1 Subsystem Configuration	200
12.4.2 Online at IPL Considerations	201
12.4.3 Switched Disconnect	202
12.4.4 APPN Minimum Switched Status	202
12.4.5 Communications Recovery	203
12.4.6 APPC Controller Description Error Recovery	204
12.4.7 Automatic Creation of APPC Controllers	205
12.4.8 Automatic Deletion of Controllers and Devices	205
12.4.9 Prestart Jobs	206
12.4.10 Client Access Mode Descriptions	207
12.4.11 Job Log Considerations	207
12.5 Testing Error Recovery	208
12.5.1 Types of Error Recovery Testing	208
12.5.2 ERP Testing Tips	209
12.5.3 Problem Determination	209
Chapter 13. Network Availability—TCP/IP Considerations	211
13.1 Saving TCP/IP Configurations	211
13.2 Saving Integrated File System Objects	212
13.3 TCP/IP Tips	213
13.3.2 Common TCP/IP Problems	215
Chapter 14. Availability Options with Hypertext Transfer Protocol	217
14.1 Backing Up HTTP Files	217
14.2 Where to Find HTTP Server Protection Setup	217
14.3 HTTP Access Control and Management	218
14.4 HTTP Problem Determination	219
14.4.1 HTTP Log File Setup Using the DDS Format	219
14.5 For More Information	220

Chapter 15. Backup and Restore for Integrated File System Objects	221
15.1 The Integrated File System	221
15.2 Save and Restore for the Integrated File System	223
15.2.1 When to Use the SAV Command	224
15.2.2 Considerations When Saving Across Multiple File Systems	224
15.2.3 Considerations when Saving Objects from the QSYS.LIB File System	226
15.2.4 Considerations when Restoring across Multiple File Systems	228
15.2.5 Considerations when Restoring Objects from the QSYS.LIB File System	229
15.2.6 Considerations when Restoring Objects to the QDLS File System	230
15.3 User-Defined File System	231
15.3.1 Mounting User-Defined File System	231
15.3.2 Saving and Restoring an Unmounted UDFS	235
15.3.3 Saving an Unmounted UDFS	236
15.3.4 Restrictions when Saving an Unmounted UDFS	236
15.3.5 Restoring an Unmounted UDFS	237
15.3.6 Restrictions when Restoring an Unmounted UDFS	237
15.3.7 Restoring an Individual Object from an Unmounted UDFS	237
15.3.8 Saving a Mounted UDFS	237
15.3.9 Restoring a Mounted UDFS	240
15.3.10 Integrated File System Commands	240
15.4 Document Library Objects	241
15.4.1 Saving DLOs	241
15.4.2 SAVDLO Enhancements	241
15.4.3 Methods of Saving Multiple Documents	242
15.4.4 DLO(*SEARCH)	242
15.4.5 Authority for SAVDLO Commands	243
15.4.6 Saving Office Services Information	243
15.4.7 Saving Mail	244
15.4.8 Saving Text Search Services Files	244
15.4.9 Restoring DLOs	244
15.4.10 Restoring New and Existing DLOs	245
15.4.11 RSTDLO Enhancements	245
15.4.12 General Performance Considerations for DLOs	245
15.4.13 Further Considerations	246
15.4.14 Authority for RSTDLO	246
15.4.15 Restoring Folders	247
15.4.16 Restoring Mail and Distribution Objects	247
15.4.17 Authority and Ownership Issues a During a Restore of DLOs	247
15.4.18 Recovery of Text Index Files for Text Search Services	248
15.5 Domino for AS/400	248
15.5.1 Why You Should Back Up a Domino for AS/400 Server	248
15.5.2 Libraries and Directories for the Domino for AS/400 Product	249
15.5.3 Backing Up the Domino for AS/400 Product	249
15.5.4 Backing Up the Domino for AS/400 Server	250
15.5.5 Backing Up Specific Dynamic Objects From Your Domino Server	251
15.5.6 Recovery of Domino for AS/400	254
15.5.7 Recovering Domino Mail	255
15.5.8 Recovering a Specific Database	256
15.5.9 Restoring Changed Objects to the Domino for AS/400 Server	256
15.6 Windows NT	258
15.6.1 Directories and Objects for Windows NT	258
15.6.2 Backing Up System Objects	259
15.6.3 Backing Up User Objects	259

15.6.4	Considerations for Back Up when Creating User Spaces	260
15.6.5	Backing Up Specific Objects From Windows NT	260
15.6.6	Restoring the Windows NT Product	262
15.7	Lotus Notes on the Integrated PC Server	264
15.7.1	Types of Storage Spaces	264
15.7.2	Backup and Recovery	267
15.7.3	Saving a Network Server Description	268
15.7.4	Saving the Server Storage Spaces	268
15.7.5	Saving the Network Server Storage Spaces	269
15.7.6	Restoring a Network Server Description	269
15.7.7	Restoring the Server Storage Spaces	269
15.7.8	Restoring the Network Server Storage Spaces	270
15.7.9	Saving by Using the ADSM OS/2 Lotus Notes Agent	270
15.7.10	Backing Up Data Using the Notes Backup Agent	271
15.7.11	Restoring Databases Using the Notes Agent	271
15.7.12	Restoring Individual Documents	272
15.8	NetWare on the Integrated PC Server	272
15.8.1	QNetWare Characteristics	272
15.8.2	Network Server Storage Spaces and Volumes	272
15.8.3	Save and Restore Overview	273
15.8.4	Types of Storage Spaces	273
15.8.5	Save and Restore Options	274
15.8.6	Saving Everything	275
15.8.7	Saving Specific Objects	275
15.8.8	Saving the Server Storage Spaces	278
15.8.9	Saving the Network Storage Spaces	279
15.8.10	Restoring Everything	282
15.8.11	Restoring Specific Objects	282
15.8.12	Restoring Server Storage Spaces	284
15.8.13	Restoring Network Storage Spaces	284
15.8.14	Other Tips and Techniques	285
15.9	OS/2 Warp Server for AS/400 (Formerly Known as LAN Server for OS/400)	287
15.9.1	OS/2 Warp Server for AS/400 Structure	288
15.9.2	Backup and Recovery for OS/2 Warp Server for AS/400	290
15.9.3	Storage Spaces	290
15.9.4	Authority Requirements for Saves	290
15.9.5	Restricted State	291
15.9.6	Tips for Saving OS/2 Warp Server for AS/400 on RISC Machines	291
15.9.7	Examples of Saving Specific OS/2 Warp Server for AS/400 Files	292
15.9.8	Restoring OS/2 Warp Server for AS/400	295
15.9.9	PC Client	297
15.10	Firewall	297
15.10.1	Saving the Firewall	298
15.10.2	Restoring the Firewall	300
15.10.3	Saving and Restoring the Filter Rules Using the COPY Command	302
Chapter 16.	Database Protection and Availability	303
16.1	Saving Database Files for Recovery	303
16.2	Logical and Physical Files in Different Libraries	304
16.3	ANZDBF Command	304
16.4	Referential Integrity Save and Restore Considerations	306
16.5	Save and Restore Tips for Trigger Programs	306
16.6	Save and Restore Relational Database Directories	307
16.7	Stored Procedures	310

16.8 Database Journaling	310
16.9 Determining Whether to Apply Journal Changes	311
16.10 Considerations for SAVCHGOBJ when Journaling is Active	312
16.11 Database Journaling Performance	312
16.11.1 Performance Tips for Journaling	313
16.12 PTFs for CHGJRN performance improvements:	314
16.13 Elimination of Lock Conflicts Between CHGJRN and RCVJRNE	315
16.14 Considerations for 1TB Maximum Access Path Size	315
16.15 Journaling of Access Paths	316
16.15.1 Access Path Journals Compared to SMAPP	316
16.16 Saving Access Paths	318
16.17 Journal Entries Considerations for V4R2	320
16.18 Journal Receiver Protection	321
16.19 Multi-member Database File Save Performance	321
16.20 Database Server Jobs	322
16.21 DB2 Multisystem	322
16.22 Restoring Distributed Files	323
16.22.1 Distributed Files Backup Considerations	323
16.23 Distributed File Back Up Scenario	324
16.24 Considerations for a Multinational Environment	325
16.25 For More Information	325
Chapter 17. Using Remote Journals to Improve Availability and Recovery	327
17.1 Remote Journal Function	327
17.2 When the Remote Journal Function Can Be Used	329
17.3 Remote Journal Transport Protocol	330
17.4 Remote Journal Replication Modes	330
17.4.1 Synchronous Delivery Mode	330
17.4.2 Asynchronous Delivery Mode	331
17.5 Performance Considerations for Remote Journal Implementation	331
17.5.1 Create Journal Receiver ASPs	331
17.5.2 Remove Internal Journal Entries	332
17.5.3 Reduce Length of Journal Entries	332
17.5.4 Check *BASE Main Storage Pool Size on Target Machine	333
17.6 Performance Considerations for Running Remote Journal	333
17.7 Remote Journal APIs	334
17.8 Remote Journal Coding Examples	334
17.8.1 Implementing Remote Journals with C	335
17.8.2 Implementing Remote Journals with RPG	339
17.9 More Information on the Remote Journal Function	350
Chapter 18. OptiConnect for OS/400	351
18.1 What OptiConnect for OS/400 Is	351
18.2 OptiConnect Hardware	352
18.3 OptiConnect for OS/400 Software Component Overview	353
18.3.1 OptiConnect for OS/400	353
18.3.2 OptiMover for OS/400 (PRPQ P84291 Product Number 5799-FWQ)	353
18.4 Bus Technology and the OptiConnect Environment	354
18.5 OptiConnect Satellites	355
18.6 OptiConnect Hub Selection	356
18.7 Dual Path OptiConnect Overview	356
18.7.1 Satellite Dual Path	356
18.7.2 Hub Dual Path	357
18.8 OptiConnect Diagrams	359
18.9 OptiConnect RPQs	359

18.10 For More Information	360
Appendix A. AS/400 Maximum Capacities	361
A.1 Limits for Database and SQL	362
A.2 Limits for Communications	365
A.3 Limits for Work Management and Security	368
A.4 Limits for Save and Restore	369
A.5 Miscellaneous Limits	370
Appendix B. Evaluating the Time to IPL	371
Appendix C. Save and Restore Rates of IBM Tape Drives for Sample Workloads	379
C.1 Comparing Performance Data	380
C.2 Lower Speed Tape Drives	381
C.3 Medium Speed Tape Drives	382
C.4 Highest Speed Tape Drives	382
C.5 Save and Restore Rates	383
C.6 Save and Restore Rates for Optical Device	385
Appendix D. OptiConnect for OS/400 Terminology and Hardware Overview	387
D.1 OptiConnect for OS/400 Terminology	387
D.1.1 An OptiConnect Cluster	387
D.1.2 Satellite and Hub Systems	387
D.2 Link and Path Redundancy	387
D.2.1 Link Redundancy	388
D.2.2 Path Redundancy	388
D.3 Hardware Overview	389
D.3.1 OptiConnect Adapter Cards and Connecting to the Network	389
Appendix E. High Availability Solutions	391
E.1 A High Availability Customer Scenario	391
E.2 When to Consider a High Availability Solution	392
E.2.1 What a High Availability Solution Is	392
E.3 DataMirror	394
E.3.1 DataMirror HA Data	395
E.3.2 ObjectMirror	395
E.3.3 SwitchOver System	395
E.3.4 OptiConnect and DataMirror	396
E.3.5 Remote Journals and DataMirror	396
E.3.6 More Information about DataMirror	396
E.4 IBM and High Availability	397
E.4.1 IBM DataPropagator Relational Capture and Apply for AS/400	397
E.4.2 DataPropagator/400 Description	398
E.4.3 DataPropagator/400 Configuration	399
E.4.4 Data Replication Process	399
E.4.5 OptiConnect and DataPropagator/400	401
E.4.6 Remote Journals and DataPropagator/400	401
E.4.7 DataPropagator/400 Implementation	401
E.4.8 More Information about DataPropagator	401
E.5 Lakeview Technology	401
E.5.1 MIMIX/400	402
E.5.2 MIMIX/Object	404
E.5.3 MIMIX/Switch	405
E.5.4 MIMIX/Monitor	405

E.5.5 MIMIX/Promoter	406
E.5.6 OptiConnect and MIMIX	406
E.5.7 More Information About Lakeview Technology	406
E.6 Vision Solutions, Inc.	406
E.6.1 OMS/400—Object Mirroring System	407
E.6.2 ODS/400—Object Distribution System	408
E.6.3 SAM/400—System Availability Monitor	408
E.6.4 High Availability Services/400	409
E.6.5 More Information About Vision Solutions, Inc.	410
Appendix F. Special Notices	411
Appendix G. Related Publications	413
G.1 International Technical Support Organization Publications	413
G.2 Redbooks on CD-ROMs	413
G.3 Other Publications	413
How to Get ITSO Redbooks	417
How IBM Employees Can Get ITSO Redbooks	417
How Customers Can Get ITSO Redbooks	418
IBM Redbook Order Form	419
List of Abbreviations	421
Index	423
ITSO Redbook Evaluation	425

Preface

This redbook serves as a companion to other sources of availability and recovery topics that most significantly includes the *Backup and Recovery* manual, SC41-5304. It offers a collection of tips and techniques from many sources, including professional consultants specializing in availability, recovery, systems management, and performance. Plus, it highlights new availability and recovery features for V4R2 and significant functions from earlier releases.

The information in this redbook assumes that you are familiar with AS/400 operating procedures, such as using commands to save your system. It also assumes that you understand basic problem determination techniques, such as where to find error messages and how to report problems.

This redbook was written for the AS/400 system administrator. System administrators are responsible for ensuring that the AS/400 system maintains a level of availability to meet business demands. Their responsibility includes:

- Ensuring the system is backed up in the event that recovery is needed
- Enforcing and monitoring the backup and recovery plan
- Planning for, implementing, and managing the appropriate hardware and software components to ensure a level of availability consistent with business requirements
- Defining system-wide values that affect save-and-restore operations, data integrity, security, and performance
- Enabling the techniques and tools to automate save-and-restore procedures
- Maintaining reliable documentation for changes made to the system

As IBM continues to enhance the availability of the AS/400 system and improve save-and-restore functions, more tools and techniques are readily accessible to maintain your system's availability. This redbook can help you understand the features of the AS/400 system, increase your system's availability, and prevent or reduce the impact of an outage.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world and is a result of a residency conducted at the International Technical Support Organization Rochester Center.

Susan Powers is an Advisory Software Engineer at the International Technical Support Organization, Rochester Center. Prior to joining the ITSO in 1997, she was an AS/400 Technical Advocate in the Support Center with a variety of communications, performance, and work management assignments. Her IBM career began as a Program Support Representative and Systems Engineer. She holds a degree in mathematics, with an emphasis in education, from St. Mary's College of Notre Dame.

Ellen Dreyer Andersen is an Advisory Systems Engineer in IBM Denmark. She has seventeen years of experience working with the System/3X and AS/400 platforms. Since 1994, Ellen has specialized in AS/400 Systems Management

with a special emphasis on performance, ADSM/400, and high availability solutions.

Rob Jones is an AS/400 Technical Advisor in Canada. He has 8 years of experience in the information technology field. He has worked at ISM, a Division of IBM Global Services, for 4 years. His areas of expertise include performance management, communications, and systems automation. He holds a degree in business from Ryerson Polytechnical Institute in Toronto.

Hubert Lye is a Software Services Specialist in Australia. He has eleven years of experience in the information technology field and has worked at IBM for eight years. His areas of expertise are AS/400 problem determination and resolution in the systems area. He holds a degree in Economics and Computing Science from Manchester Polytechnic in the United Kingdom.

Petri Nuutinen is a Systems Support Engineer in Finland. He has fifteen years of experience, starting with the S/38 and has been working with the AS/400 since early 1987. His main areas of expertise are AS/400 problem determination, work management, and performance tuning. He holds a degree from the University of Finland.

Thanks to the following people for their invaluable contributions and advice in the production of this document:

International Technical Support Organization, Rochester Center

Marcela Adan
Jim Cook
Jarek Miszczyk
Suehiro Sakai
Fant Steele

IBM Rochester Laboratory

Bernie Begin
Jim Bonalumi
Pamela Bowen
Mike Denney
Mark Diez
Jim Flanagan
Bob Gintowt
John Halda
Don Halley
Steve Hank
Jelan Heidelberg
Allan Johnson
Paul Koeller
Kathy Mack
Dawn May
Scott Maxson
Pundi Madhavan
Jerry Miller
Craig Nordmoe
Dave Novey
Dick Odel
Ron Peterson
Luz Rink
Joe Rizzo

Debbie Saugen
Edward Stavana
Chuck Stupca
Jeff Tenner
Marty Thompson
R. J. Traff
Judy Trousdell
Tony Tschida
Jeff Vettel
Larry Youngren
Paul Wolf

IBM UK Technical Support

Paul Kirkdale

IBM Technology Solutions Center

Selwyn Dickey
Brenda Thompson

Endicott Software Development Lab

Mark Bullock
Joseph Caldwell
John Hall
Susan Hall
Rich Hock
John Martz
Joseph Miller
Frank Paxhia

AS/400 Support Line

Luis Barajas
James Hall
Richard Halleen
Mike Moiwood
Kent Morris
Bill Osler
Sue St. George
Peter Schmitt

IBM AS/400 Competency Center

Sue Baker
Eric Hess

Partners in Development

Amit Dave
Kent Milligan

AS/400 Skills Development

Olga Saldivar
Michael Cameron-Smith

Representatives from the Large AS/400 User Group (LUG)

Jerry E. Linde, Communications Data Services, Incorporated
Rod Flinn, Circuit Cities Stores, Incorporated
Don Birch, Nintendo of America

Robert J. Cargill, Oriental Trading Company, Incorporated
Gary S. Lagarde, Reynolds Metals Company

Representatives from IBM Business Partner Organizations

DataMirror Corporation—Wayne Nathanson
Lakeview Technology
Vision Solutions, Inc.—Fred Grunewald

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 425 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:
For Internet users <http://www.redbooks.ibm.com/>
For IBM Intranet users <http://w3.itso.ibm.com/>
- Send us a note at the following address:
 redbook@us.ibm.com

Chapter 1. Introduction to AS/400 Availability and Recovery

The AS/400 system maintains a strong reputation for reliability. AS/400 users trust it to run their most critical applications. Over the past two years, these users have averaged less than nine hours of unplanned downtime per year, according to data reported by IBM. The data also indicates that a single AS/400 system delivers an average of 99.9% or more availability.¹

System availability is critical to your organization's efficiency and effectiveness. Availability is determined by system design, configuration, management, and level of user control. It is calculated by the number of hours in a year (8 760) that the system is available, minus the number of hours of unplanned downtime. Most importantly, system availability depends on the preparation and testing of a complete availability and recovery plan.

This redbook serves as a source of AS/400 availability and recovery information. It includes details about available methods and options to help you reach your required level of system availability. It also describes techniques that you can incorporate into your plan to maximize system availability and ease of recovery as you need it.

Topics addressed in this redbook include:

- Hardware availability solutions
- Procedures to effectively save and restore the AS/400 system
- Selecting a suitable tape drive for your backup
- Using remote journals and database protection methods
- Automating availability
- Managing a mixed release environment
- Licensed program considerations
- Availability and communications facilities
- PTFs and availability
- Minimizing the time to IPL
- Protecting the integrated file system environment
- OptiConnect for availability
- High availability solutions
- Limits to growth

Throughout this redbook, there are many references to the manual *Backup and Recovery*, SC41-5304. This redbook was designed as a companion to this manual and other related backup and recovery publications. Together, these guides can help you maintain the reliability of your AS/400 system and establish recovery during an unplanned or planned outage.

To help you understand availability and recovery as it is presented in this redbook, the next section offers a brief review of these features since the introduction of the AS/400 in 1988. Following this section is a list of commonly asked questions regarding system availability and recovery.

¹ The measurement is based upon feedback from:

- IBM's Field Data Management System Hardware Service Call Reports for 1996, 1997, and 1998
- "Customer Survey of AS/400 High Impact Outages," IBM, July 7, 1996
- System and availability information reported to IBM by the Large User Group for 1997 and 1998

1.1 A Historical Perspective of AS/400 Availability Enhancements

Since IBM introduced the AS/400 system in 1988, many enhancements have been made to make the system more readily available and easier for users to recover upon any given failure. Historically, the main availability enhancements tracked by AS/400 software releases are:

Version 1 (including V1R1, V1R2, and V1R3): Disk mirroring support was added to protect against DASD failures. Users could protect DASD devices concurrently without bringing down the system. IBM 3490 tape technology was also introduced, which provided increased capacity for archives, saves, and unattended backups.

Version 2 (including V2R1, V2R2, V2R3, and V3R05): An integrated battery power unit was introduced on the high-end models to protect against power outages. The new Save While Active function reduced the downtime to perform save operations. Device Parity Protection (RAID) support was added to provide more cost-effective protection against DASD failures. Plus, tape capacities and performance continued to increase.

Version 3 (including V3R1, V3R2, V3R6, and V3R7): System Managed Access Path Protection (SMAPP) was introduced to control and reduce initial program load (IPL) time required to rebuild access paths after a failure. Faster IPLs and the elimination of the need to IPL for regenerating temporary addresses appeared in 64-bit RISC machines. The addition of IBM 3590 tape technology, faster I/O busses on RISC machines, and support for a larger tape block size resulted in vastly improved save and restore performance and capacity. Backup Recovery and Media Services (BRMS/400) was added to support automated and policy-driven archive, backup, recovery, and management of tape media. Plus, DASD devices could now be concurrently added on RISC systems without bringing down the system. And, many PTFs could be installed or applied concurrently without requiring an IPL.

Version 4 (which includes V4R1 and V4R2): System availability increased, with a reduction of up to 50% in the time needed to complete an IPL. Scheduled dedicated system time for applying cumulative PTF packages was reduced by up to one-third. New support was added on the save commands for generic library names and omit options to provide maximum flexibility for defining which libraries and objects to save. This support made it easier to use multiple tape drives to further reduce the time that the system is unavailable during save operations. Multiple concurrent save or restore operations with two or more tape units could save or restore objects to a single library, or save or restore DLOs into a single ASP. The most common locking conflicts, while an object is saved with the Save While Active function, were eliminated. System and network availability was increased due to enhancements in starting and stopping APPC communications and streamlining the logging and recovery of errors. A remote journal function was added to replicate journal entries from a local system to journals and journal receivers located on a remote system for hot site backup, data replication, and high availability applications.

The following table offers a detailed summary of availability enhancements by release.

<i>Table 1 (Page 1 of 3). AS/400 Availability Enhancements by Release</i>	
VnRn	Availability Function
V1R1 and V1R2	<ol style="list-style-type: none"> 1. Checksum protection 2. 3490 tape drive with autoloader capacity for up to 3.6GB unattended backup
V1R3	<ol style="list-style-type: none"> 1. DASD mirroring 2. Concurrent maintenance of DASD 3. Delayed power down 4. Predictive analysis on 9335 DASD 5. DASD attention SRCs 6. Enhanced ASP 7. Save and restore performance improvements 8. 3490 tape drive model allowing 14.4GB unattended save 9. 9336 DASD
V2R1	<ol style="list-style-type: none"> 1. Checksum protection by ASP 2. Battery backup in high-end models 3. Duplicate tape function 4. 3490 tape compression capabilities providing for more than 62GB on an unattended save 5. 8 MM tape capability providing 2.3GB saves 6. Faster one-fourth inch cartridge tape drive 7. IPL time improved 8. Restart-tape processing at checkpoints 9. RCLSTG time improvement 10. SAVCHGOBJ performance improved
V2R2	<ol style="list-style-type: none"> 1. ASP management improved 2. Save While-active function 3. Saving office support enhanced 4. 3490 Model E tape drive 5. Unattended save with Operational Assistant feature 6. 9337 disk array with RAID technology and redundant power
V2R3	<ol style="list-style-type: none"> 1. Save While Active across multiple libraries 2. Automated operations enhanced 3. Save and restore performance improved to allow additional functions in a non-dedicated environment 4. Migration capabilities to RAID without reloading 5. 3490 on 9404 systems 6. Operating system RAID support for 9406 systems
V3R05	<ol style="list-style-type: none"> 1. Redundant N+1 power supplies and regulators on RISC hardware
V3R1	<ol style="list-style-type: none"> 1. SMAPP 2. QDBSRVnn database cross-reference files and jobs 3. Reset and reload of an IOP without an IPL 4. Parallel journal apply 5. System managed change journal support 6. Parallel database IPL recovery 7. Faster small object restore 8. BRMS and SystemView products 9. 3590 tape drive 10. Immediate PTFs
V3R2	<ol style="list-style-type: none"> 1. Ability to CHGJRN when a RCVJRNE operation is active

Table 1 (Page 2 of 3). AS/400 Availability Enhancements by Release

VnRn	Availability Function
V3R6	<ol style="list-style-type: none"> 1. The WRKASP utility for user-friendly access to ASPs 2. 30% reduction in normal IPL times 3. Elimination of IPLs to regenerate addresses 4. 67% increase in maximum save rate up to 30GB per hour with 3590 5. Allow more than nine concurrent save and restore jobs 6. Continuously powered main store 7. ObjectConnect commands 8. Dynamic priority scheduling 9. Concurrent service of internal tape 10. Automatic ECS reporting when CPU is down 11. Remote access to control panel when CPU is down 12. Ability to CHGJRN when RCVJRNE is active 13. Concurrent add of DASD
V3R7	<ol style="list-style-type: none"> 1. Performance and reliability for device recovery 2. Mirror DASD on a MFIOp to another IOp 3. Increase SAVDLO and RSTDLO limit to two million DLOs 4. Concurrent SAVDLO and RSTDLO for different ASPs 5. Communications error recovery improvements 6. R/DARS 7. 57% increase in maximum save rate up to 47GB per hour 8. Twenty times the number of logical and physical files in a database network 9. Multiple concurrent SAVDLO and RSTDLO against different ASPs 10. Four times the improvement in main store copy to disk 11. More prestart jobs for PC connect time servers 12. Bus-level concurrent repair of IOps

Table 1 (Page 3 of 3). AS/400 Availability Enhancements by Release

VnRn	Availability Function
V4R1	<ol style="list-style-type: none"> 1. The ability to select or omit certain functions for the RCLSTG process 2. Optimization for additional save commands using the USEOPTBLK parameter 3. Up to 50% reduction in IPL time 4. Main store dump time improved 5. Main store dump diagnostics available earlier in IPL 6. Apply LIC PTFs with a single IPL 7. Eliminate most save-while-active restrictions after checkpoint 8. Alternate IPL allowed from any I/O bus 9. Multiple concurrent SAVOBJ commands against a single library 10. Multiple concurrent SAVDLO commands against a single ASP 11. Generic folder names on SAVDLO and new OMITFLR parameter 12. Generic OMITLIB values to break up saves across multiple devices 13. Options to omit configuration and security data for SAVSYS 14. Improved save and restore performance especially for multiple member files 15. Decreased time for save-while-active checkpoints 16. RSTAUT performance improvements 17. Operator control of an IPL 18. Communications error recovery improvements 19. Licensed program install performance improvements 20. RAID MFIOPs 21. Concurrent DASD repair on 9402 systems 22. Concurrent maintenance of base power supplies and cooling fans for some models
V4R2	<ol style="list-style-type: none"> 1. Remote journal support 2. Multiple concurrent RSTOBJ commands against a single library 3. Multiple concurrent RSTDLO commands against a single ASP 4. Generic LIB values to save groups of libraries with similar names 5. Expanded OMITLIB and new OMITOBJ values 6. New prompts for SAVE menu options (Vary off IPCS, Unmount user-defined file systems, PRTSYSINF) 7. Save While Active to any valid target release 8. Reduction in abnormal IPL time 9. Up to 30% less time to perform ENDSBS, ENDSYS and PWRDWNSYS 10. Communications error recovery improvements

1.2 Frequently Asked Questions about Availability and Recovery

This section serves as a reference to solutions for commonly asked questions about availability and recovery. The questions are divided into six categories, which include:

- Disk management
- Database management
- Saving or restoring the system
- System availability factors
- System management
- Communications or network issues

You can quickly locate the information you need without scanning the entire redbook. In some cases, one question leads to more than one topic, and many questions point to the same chapter.

1.2.1 Disk Management

The following questions and answers relate to disk management:

- **I want to implement user ASPs, but they look difficult to manage. Is there a tool I can use?**

Yes. The WRKASP utility provides a user-friendly menu interface to manage user ASPs. Refer to Section 9.4, “WRKASP Utility” on page 120.

- **What options are available to protect the system in the event the load source unit has an unrecoverable error?**

See Section 3.1, “Load Source Protection” on page 23.

1.2.2 Database Management

The following questions and answers apply to database management:

- **I heard about SMAPP, but turned it off. Is that a good idea?**

No. Please look at Section 4.4, “System Managed Access Path Protection” on page 38, to see what is new for SMAPP.

- **It takes a long time to rebuild. Is there a way I can speed it up without doing a manual IPL?**

If you saved the data with the access paths on a tape, the answer is yes. Put the rebuild of access paths on hold and restore the file from tape. See Chapter 16, “Database Protection and Availability” on page 303, for further considerations about access paths.

- **How do I obtain a better understanding of my database structure and the location of the physical and logical files?**

Turn to Section 16.9, “Determining Whether to Apply Journal Changes” on page 311.

- **If I have DASD mirroring, do I need journals?**

Yes. Mirroring is designed to protect from loss of a disk drive by keeping the system operating while the drive is repaired or replaced. Mirroring cannot protect you from data corruption or user error or object damage caused by an abnormal system end. This is what journals are designed for. Refer to Section 16.8, “Database Journaling” on page 310.

- **I heard that journals can be located on a system separate from the associated database. Can this help me offload my production system? Where can I find out more?**

Yes. Remote journal support capabilities within OS/400 provide enablers that allow for data changes to be synchronously recorded on two or more systems. The speed and integrity of the high availability application is enhanced and in-flight data is protected across a failure. Refer to Chapter 17, “Using Remote Journals to Improve Availability and Recovery” on page 327, and Appendix E, “High Availability Solutions” on page 391, for information on remote journals and high availability solutions.

- **I am considering the development of journal replication in my application through the remote journal function to free up some resources on my production machine. How can I test the new APIs to determine what kind of performance I can expect?**

The APIs used for remote journal function are listed in Section 17.7, “Remote Journal APIs” on page 334. See Section 17.8, “Remote Journal Coding Examples” on page 334 to see how to create a program with remote journal interfaces.

1.2.3 Saving or Restoring the System

This section presents frequently asked questions about saving or restoring the system:

- **How can I make sure that I can restore everything that I saved back to my system?**

Test your recovery. Refer to Section 5.7.2, “Additional Considerations” on page 61 for details.

- **What happens when I do not use the SAV command?**

If you do not use the SAV command, you only save the root directory (QSYS.LIB) using the SAVSYS command. All other directories, such as QDLS and QLANSRV, are not saved. Refer to Section 15.1, “The Integrated File System” on page 221, in this redbook for a discussion on this topic.

- **My SAVSYS operation takes too long. How can I reduce the time that my system is in a restricted state?**

See Section 5.3, “Omitting Objects on a SAVSYS Operation” on page 53.

- **Saving my production database takes too long. How can I reduce the elapsed save time of my production libraries?**

Through hardware and software enhancements you can run multiple parallel backups. These backups can execute concurrently. You can divide your database to send equal amounts of data to each of your tape drives. For an example of how this can be done, see Section 5.4, “Concurrent Save Operations” on page 55.

- **I am worried that I do not save my database in the best way to ensure a fast restore and recovery. Any recommendations?**

Yes, turn to Appendix C, “Save and Restore Rates of IBM Tape Drives for Sample Workloads” on page 379, for a description of the various ways to save a database with a fast recovery in mind.

- **Can I restore an individual object saved with a SAVLIB (*NONSYS, *IBM, or *ALLUSR) command?**

Yes. First, display the tape using the Display Tape (DSPTAP) command, and indicate *SAVRST for the data type parameter. Find the object you need to restore and note the file label ID. On the Restore Object (RSTOBJ) command, specify the name of the object and the library from which it was saved. For the RSTOBJ label parameter value, use the file label ID you noted from the DSPTAP command. Refer to Chapter 5, “Save and Restore for Availability and Recovery” on page 49 for more save and restore considerations.

- **Can I restore subsystem descriptions, QHST files, job queues, or commands from a SAVSYS tape?**

Yes. First, display the tape using the Display Tape (DSPTAP) command, and enter *SAVRST for the data type. Find the object that you need to restore, and note the file label ID, which you need for the Restore Object (RSTOBJ) command. On the RSTOBJ command, enter the name of the object and the library from which it was saved. Use the file label ID that you noted from the DSPTAP command for the RSTOBJ label parameter value.

- **Why are my private authorities gone when I delete a library and restore from backup?**

Private authorities are not saved with the objects. To recover the private authorities of objects you have already restored, you must restore the user profiles and run the Restore Authority (RSTAUT) command. Refer to Chapter 5, “Save and Restore for Availability and Recovery” on page 49, for more save and restore considerations.

- **Do I have to implement BRMS to back up spool files?**

No. Refer to Section 5.10, “Save and Restore Spooled Files” on page 69, for some sample programs.

- **Is there a way to move objects between systems without using a tape drive if my SNADS connection is down?**

Yes. On your system you have a set of commands ready to be used. Refer to Section 5.9, “ObjectConnect/400” on page 65.

- **How do I save Volume SYS in NetWare? Do I need to vary off the Integrated PC Server?**

Section 15.8, “NetWare on the Integrated PC Server” on page 272 describes how to manage the NetWare file system.

- **What tips can you offer on saving NetWare objects?**

See Section 15.8.7, “Saving Specific Objects” on page 275, for more information on tips for saving NetWare objects. Also look at Table 32 on page 273, for a summary of QNetWare objects.

- **I want to make sure that my PC users’ data is backed up regularly. How can I do that?**

See Section 9.1, “ADSTAR Distributed Storage Manager/400 (ADSM/400)” on page 109.

- **After I upgraded to V4R2, my first save used more tape storage than previous saves. What happened?**

For V4R2, database files are converted to lengthen the current header extension area by 512 bytes of storage per file. This conversion takes place when the object is first used. The additional storage affects save media. For

example, 10 000 files saved to tape require an additional 5 120 000 bytes in V4R2. (10 000 X 512 = 5 120 000).

1.2.4 System Availability Factors

This section provides answers to commonly asked questions about system availability:

- **How can I reduce my IPL time?**

See Sections 4.5, “Changing IPL Attributes” on page 42 and 4.4, “System Managed Access Path Protection” on page 38.

- **How can I automate system tuning?**

You can let the system do it for you. See how this is done in Section 8.3.1, “Automatic Performance Adjustments—It Is Worth Another Look” on page 100.

- **Can I affect the way the performance adjuster works?**

Yes. There are several values that you can alter to affect the performance adjuster. In addition, algorithms for the performance adjuster changed since it was first introduced to make it more efficient. See Section 8.3.3, “Tuning the System Tuner” on page 104, for more information.

- **A few hours after an IPL, I sometimes have mysterious “slowdowns” on my system. Why does this happen?**

Take a look at Section 10.2, “Work Control Block Table” on page 141, for a few suggestions.

- **What levels of availability can I achieve on the AS/400 system?**

There are several levels of availability for which you can plan. Refer to Section 2.3, “Levels of Availability” on page 17, Section 5.4, “Concurrent Save Operations” on page 55, Section 5.7, “Save and Restore Scenario” on page 58, and Section 5.6, “Use Optimum Blocking for Save and Restore” on page 57.

- **I heard that the AS/400 system can be clustered. What does that mean?**

Clustering is a group of separate computers that behave as if they were a single system. A workstation interacts with a cluster as if it were a single highly reliable server.

AS/400 clusters are enabled by database replication and fail over support provided by high availability business partner applications. In the event of a planned or unplanned outage, operations can be quickly restarted on a backup system containing a replicated copy of the system, applications, and data.

- **I heard that OptiConnect/400 can be used to create a clustered AS/400 solution. How does the Opti/Connect400 software work and how are the machines connected? How does OptiConnect/400 provide a high availability solution?**

IBM’s OptiConnect hardware and software components enable a user to customize a high availability solution. Refer to Chapter 18, “OptiConnect for OS/400” on page 351.

- **What does the expression “a high availability solution” mean?**

High availability is the characteristic of a system that delivers an acceptable or agreed-upon level of service during scheduled periods of operation. High

availability systems recover from failures of major hardware components without a loss of data or loss of time to restore data. Think of high availability as a tool that allows you to keep your system up and running. Please refer to Appendix E, “High Availability Solutions” on page 391, for more information on high availability software solutions.

- **Do I have to purchase a separate product to obtain a high availability solution?**

No. High availability solutions can be written by application programmers or you can purchase a separate high availability product. See Chapter 18, “OptiConnect for OS/400” on page 351, and Appendix E, “High Availability Solutions” on page 391.

OptiConnect/400 hardware and software offer a variety of options in developing clustered systems and hot site backups of your existing environment. To see what hardware is required and how the environments are connected, see Chapter 18, “OptiConnect for OS/400” on page 351, and Appendix D, “OptiConnect for OS/400 Terminology and Hardware Overview” on page 387.

- **We must maintain continuous operations and never close down the system for planned or unplanned outages. What do you recommend?**

Look at a high availability solution involving two or more systems. See Appendix E, “High Availability Solutions” on page 391, for more information.

- **Can exceeding storage limits affect system availability?**

Yes. Refer to Appendix A, “AS/400 Maximum Capacities” on page 361 for a table of maximum capacities for V4R2 AS/400 systems.

1.2.5 System Management

The following questions and answers apply to system management:

- **Can tapes created using the SAVSYS command be used for installing software?**

Tapes that are created using SAVSYS are meant for recovery operations. Therefore, you *cannot* use them for an automatic installation process. SAVSYS tapes do not provide a full backup of your system. See Chapter 5, “Save and Restore for Availability and Recovery” on page 49 and Section 9.5, “Print System Information Tool” on page 123.

- **Can I restore an object to a lower release system?**

Yes. You must specify the target release on the save command. Refer to Chapter 6, “Save and Restore Considerations for Mixed Release Environments” on page 75, for more information on saving and restoring objects in a mixed release environment.

- **I delay upgrading software levels because I am afraid I will lose my configuration. Is there a way to prevent this?**

Yes. Use the PRTSYSINF command to gather necessary documentation on how your system is configured, both before and after any upgrade. Refer to Section 9.5, “Print System Information Tool” on page 123 for more information.

- **I need to save data from my V4R1 system to another system that is running at V3R6. Can I do that?**

Look at Chapter 6, “Save and Restore Considerations for Mixed Release Environments” on page 75, for a description of considerations for

maintaining a multiple release level environment and for a list of valid parameters on the various save commands.

- **Can I distribute software packages, install them on my remote systems, and perform an IPL if necessary?**

Yes, you can. Please refer to Section 9.9, “IBM SystemView Managed System Services for AS/400” on page 130, for details.

- **I want to distribute software packages to my remote AS/400 systems, as well as to non-AS/400 systems. What is the best way to do this?**

See Section 9.8, “IBM SystemView System Manager for AS/400” on page 128, for a solution.

- **I delay managing PTFs because it is difficult to do and requires a lot of time. Has this process been improved?**

Actually, that is not required. See Section 11.3, “Applying and Removing PTFs and IPLs” on page 170, for details.

- **Can I use my Internet access to read Preventive Service Planning information?**

Yes. See Section 11.2, “Preventive Service Planning” on page 169.

- **Can I download and apply PTFs for a later release before installing the later release?**

Yes. See Section 11.7, “Applying PTFs for the Next Release Prior to Installing the Next Release” on page 177, to learn how this is done.

- **Why were PTFs not applied during last night’s IPL using the Job Scheduler?**

The Job Scheduler can only perform IPLs to the B side. To apply Licensed Internal Code (LIC) PTFs, the system needs to go to the A side. See Section 11.3.1, “LIC PTF Apply and System Availability” on page 171, for more information.

- **I loaded and applied a PTF on my system and performed an IPL, but the PTF is still not activated. Can I verify that all required PTFs are on my AS/400 system before performing additional IPLs?**

Yes. See Section 11.5, “Conditional PTFs” on page 175, to learn how the PTF process has changed.

- **Can the SNDPTF command send PTFs to other systems, or do I need the Managed System Services for AS/400 product?**

Yes. The SNDPTF command sends the PTFs when Managed System Services is not installed on the service requestors system. However, the *Extent of Change* parameter is not functional. For example, if *APY(*TEMP)*, the PTF is sent but not applied.

See Section 11.8, “Distributing PTFs” on page 177, for more information on distributing PTFs, and see *System Manager Use*, SC41-5321-01, for information on distributing PTFs and setup considerations.

- **I want to automate my batch backup jobs process. How do I start?**

You can do this with the OS/400 job scheduler as described in Section 9.6, “OS/400 Job Scheduler (Part of the Operating System)” on page 125.

- **I want to automate the interdependent batch backup jobs process. Is this possible?**

Yes, with Job Scheduler/400. Look at Section 9.7, “IBM Job Scheduler for OS/400” on page 127.

- **How can I keep track of what has been backed up from my system and on which media it is stored?**

See Section 9.3, “Backup, Recovery, and Media Services for AS/400 (BRMS)” on page 114.

- **Do I need BRMS/400 to manage my ADSM/400 generated tapes?**

No, ADSM/400 can do that on its own. Refer to Section 9.1.5, “Total Disaster Recovery with ADSM/400 Version 2” on page 112, for an explanation.

- **I need to perform a reclaim storage on my system, but I am worried that it will take too long. Can I figure out how long it is going to take?**

You can estimate the time based on a number of factors. See Sections 10.10, “Reclaim Storage” on page 153, and 10.10.5, “Completion Time for Reclaim Storage” on page 157.

- **What does the RCLSTG operation do while it is running?**

The RCLSTG process issues messages. Refer to Section 10.10.3, “Reclaim Storage Status Messages” on page 156.

- **My AS/400 system is sometimes heavily loaded with users running queries and various statistics. Is there a way to offload it?**

Yes, with a second AS/400 system and a data replication tool, you may let your read-only applications run on a second machine. Interactive users are not affected by large queries or other heavy read-only jobs. See Appendix E, “High Availability Solutions” on page 391, for details.

1.2.6 Communications or Network Issues

This section addresses questions about communications or network issues:

- **Is there a way to audit the use of HTTP accesses?**

Yes. You can query the log files. Refer to Section 14.3, “HTTP Access Control and Management” on page 218.

- **Since I upgraded to V4, the QCMN routing entries no longer exist. prior to V4, I set up security. How do I create the same setup through the QPASTHR servers?**

You can do this by using the QRMTSIGN program. See *Remote Work Station Support*, SC41-5402.

Chapter 2. Availability and Recovery Concepts

Imagine a situation that brings your computer operations down for an entire day or week. Imagine losing all company system stored data, all on-site backup tapes and the system process. If you ask the question “What now?,” it is already too late. The only effective way to cope with computer operation failures is to have a comprehensive, fully tested backup recovery solution in place *before* you need it. This redbook helps you plan for an efficient backup and recovery solution.

The term failure, as mentioned in this book, refers to an interruption of the information processing services that cannot be corrected within an acceptable predetermined time frame. Obvious examples are floods, tornados, fires, and explosions. A breakdown of the frequency of these “hard” disasters is shown in Figure 1.

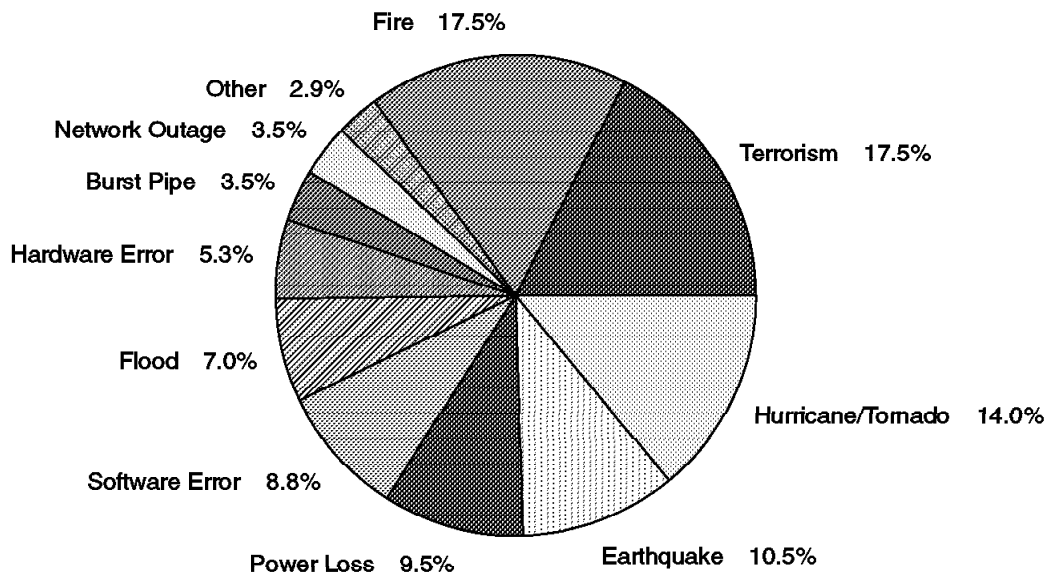


Figure 1. Disaster Frequency by Type

Sometimes, however, a minor problem that appears to be recoverable through normal problem management procedures can escalate into a disaster, such as corruption in a database. You recover the data from a backup tape, but then have further data errors. If the source of the problem cannot be identified and resolved, and the problem escalates throughout the day, management may consider this to be a disaster situation.

This chapter introduces the concepts of disaster recovery and stresses the importance of a reliable and tested backup recovery plan. Once the plan is realized, your system’s level of availability is improved.

2.1 The Importance of Backup

Organizations depend on information systems to help manage their business and to stay competitive. Most businesses require a high, if not continuous, level of information system availability—systems that are disaster tolerant. A lengthy outage of computing information can result in a significant financial loss, including the potential loss of customer credibility and subsequent market share.

Management must determine the time frame that moves an outage from a “problem” to a “disaster” status. Most organizations accomplish this by performing a *business impact analysis* to determine the maximum allowable downtime for critical business functions. One publication you can use to understand the implications of downtime on your business is the redbook *Fire in the Computer Room, What Now? Disaster Recovery: Planning for Business Survival*, SG24-4211.

The ability to successfully recover from a failure within a predetermined time frame is a critical element of an organizations’ strategic business plan. A recovery plan can help you avoid the negative financial impact of losing information, the technology, or the facilities. The University of Minnesota revealed this by conducting a study of information technology outages in 60 Minnesota enterprises representing the banking, commercial, industrial, and insurance industries. The outages ranged from two days to 5.6 days. After recovery from the disaster, the study showed that:

- Twenty-five percent of the enterprises went bankrupt.
- Forty percent were bankrupt within two years.
- Less than seven percent were in business after five years.

Recommendation

View disaster recovery planning as an insurance policy for your business.

Manual procedures, if they exist at all, are only practical for a short period of time.

2.2 Backup and Recovery Planning

The main activities required in planning and implementing a recovery plan are shown in Figure 2 on page 15. Use the steps following the table as a guideline to help you develop your backup and recovery plan.

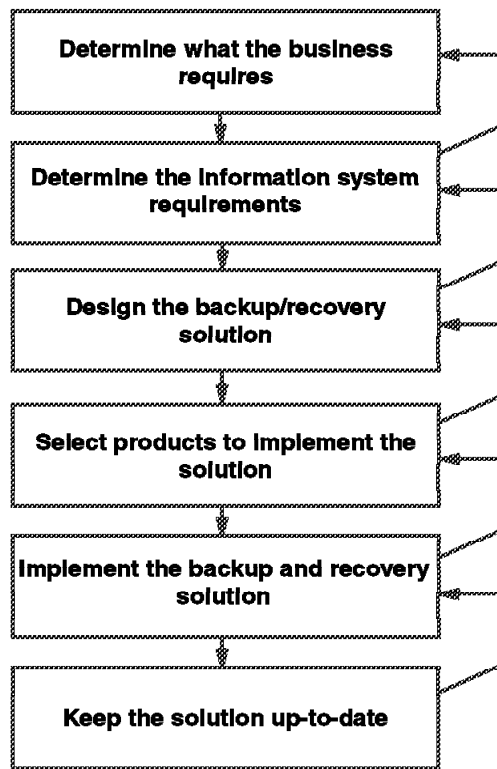


Figure 2. A Structured Approach to Disaster Recovery

1. *Determine what the business requires*

Many business processes depend completely on information systems that the business cannot effectively perform in the event of an information system failure. Business processes must be evaluated for their negative impact, that is, loss of business and revenue. The evaluation reveals what your business priorities are and the time scale that is required for the recovery for each business process.

2. *Determine the information system requirements*

Once the business requirements are established, convert them into information systems terms, or into a context that the system administrator can use. The result should be a matrix showing (for each application) the required recovery time, maximum allowable data loss, computing power required to run, disk storage required, and dependencies on other applications or data. Cooperation across organization and departmental boundaries is necessary to reach a consensus on requirements.

3. *Design the backup and recovery solution*

Describe, in a generic way, any special hardware and software functions required for the backup and recovery processes, the recovery configuration, the network and interconnection structures, and the recovery location. A high level design may be sufficient to estimate the cost of the solution. If the cost is too high, the solution needs to be reworked until both the cost and the solution are agreed upon.

4. *Select procedures and products to implement the solution*

Develop processes and products that can be used together to support the recovery design solution you developed. Products consists of hardware, software, and possibly the selection of an alternate system or site. Processes are the documented steps for all aspects for the backup and recovery plan.

5. *Implement the backup and recovery solution*

The development of the backup and recovery plan requires cooperation with many departments and sponsorship by management. Set up the recovery team to develop and implement backup and recovery procedures. Document the recovery steps to execute the recovery plan for each type of failure.

6. *Keep the solution up-to-date*

Put procedures in place to ensure that the backup and recovery solution remains viable. Develop and implement procedures for the maintenance, testing, and auditing of the backup and recovery plan. The success of the plan is only as good as its evaluation.

Four factors, as depicted in Figure 3, that you need to consider in any solution are:

1. How fast must the recovery be accomplished?
2. How much data can be lost?
3. What type of disaster does the solution cover?
4. What is the cost of the solution?




Figure 3. *Outage Cost and Outage Time*

Be aware that some requirements are trade-offs, or are mutually exclusive. The more stringent the requirements, the higher the cost of the solution.

Balance these factors to develop a solution which best meets the business needs at an acceptable cost.

Note: The backup and recovery plan forms the basis for a disaster recovery plan for information systems. Though the planning is likely to be executed by information system professionals, the recovery plan should be owned and driven by executive management.

Use the redbook *Fire in the Computer Room, What Now? Disaster Recovery: Planning for Business Survival*, SG24-4211, to help build a recovery plan.

2.3 Levels of Availability

Systems experience both planned and unplanned outages. Systems can be classified according to the degree to which they cope with different types of outages. These levels include:

- Continuously operational:
Systems that reduce or eliminate planned outages
- Highly available:
Systems that reduce or eliminate unplanned outages
- Continuously available:
Systems that reduce or eliminate both planned and unplanned outages

Your system is as available as it is planned or designed to be. Orient your implementation choices toward your desired level of availability. These levels include:

- High availability:
Systems provide high availability by delivering an acceptable or agreed upon level of service during scheduled periods of operation. The system is protected in this high availability type of environment to recover from failures of major hardware components like CPU, disks, and power supplies when an unplanned outage occurs.
- Continuous operations:
A continuous operation system is capable of operating 24 hours a day, 365 days a year with no scheduled outage. This does not imply that the system is highly available. An application can run 24 hours a day, 7 days a week yet be available only 95% of the time because of unscheduled outages. When unscheduled outages occur, they are typically short in duration and recovery actions are unnecessary or minimal.
- Continuous availability:
This type of availability is similar to continuous operations. Continuous availability systems deliver an acceptable or agreed upon service 7 days a week 24 hours a day. They add to availability provided by fault tolerant systems by tolerating both planned and unplanned outages. Continuous availability must be implemented on both a system and application level. By doing so, you can avoid losing transactions. End users need not be aware a failure or outage has occurred in the computing environment.

Note: The computing environment consists of the computer, the network, and all workstations.

The levels of protection that a system offers depends on hardware, software, and application components. Many components available for the AS/400 system are described later in this redbook. Figure 4 on page 18 shows some of the hardware and software options available for your AS/400 system.

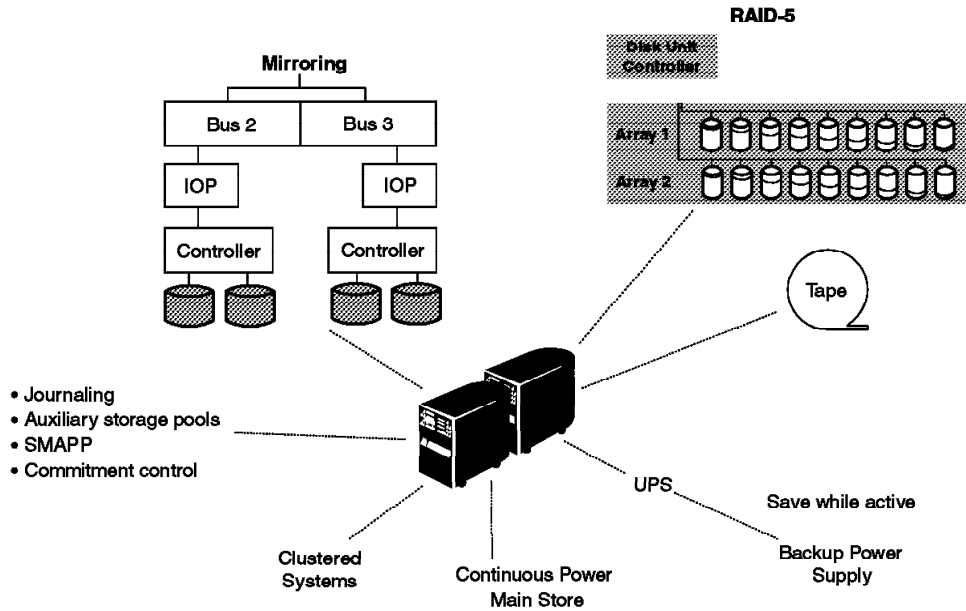


Figure 4. AS/400 Availability Components

2.4 Types of Unplanned Outages

Unplanned outages can result from any of several potential failure types:

- Power loss
- DASD loss
- System failure
- Program or human error
- Site loss
- Performance degradation

The AS/400 system has multiple solutions to address the problems associated with each of these failure types. This redbook discusses many, but not all, of the solutions available.

Note: Performance management is not viewed by some as an availability option. However, response time for interactive and batch jobs affects the *availability of information* in a timely manner. It is from this perspective that this redbook addresses performance.

See Chapter 8, "Save, Restore, and System Performance for Availability" on page 95, and Appendix C, "Save and Restore Rates of IBM Tape Drives for Sample Workloads" on page 379, for a further discussion of availability related performance topics.

2.4.1 Power Failure

The loss of room, floor, building, or site utility power can cause major disruption for jobs on the system. This creates the potential for a long IPL and the chance for damaged objects. Recovery time varies considerably based on the AS/400 model (CISC or RISC), the protection options installed, the type of applications, database design, activity on the system at the time of failure, and (primarily) AS/400 availability options in place at the time of power failure.

2.4.2 Unprotected DASD or Disk Failure

Disk failures are when one or more unprotected disk units in an ASP fail. Disk failures are also called DASD, disk arm, disk actuator, or disk spindle failures. Recovery time varies considerably (from none to a lot) and can require a reload of the entire ASP where the failed disk unit resides. Recovery time varies depending on performance of the AS/400 model, types and numbers of disks, size of data to be restored, save strategy, number and speed of tape units used to restore, the type of objects to be restored, and the method used to protect the disk drives against failure.

2.4.3 System Failure

System failures are caused by the loss of system hardware and software failures such as a memory card, CPU, disk controller, cables, program defects, and so on. Recovery time ranges from none to some and may be similar to a power failure. Recovery time is based on the AS/400 model (CISC or RISC), options installed, type of applications, database design, activity on the system, and primarily the AS/400 availability options in place at time of system failure.

2.4.4 Human Error or Program Failure

A human or program error is usually the result of a program or person that incorrectly executes and deletes or updates records or a file in an undesirable fashion. Typically this does not impact the entire system and business can continue. How severe the impact is depends on the application affected. For example, a mission critical order entry application outage may be more severe than the unavailability of an inventory application. Recovery time varies widely from some to a lot and depends on the number of humans or programs involved, the ability to realize that the problem is occurring, and the ability to debug or repair the program (or humans). Recovery can include restoring the database files that are affected, so some of the issues noted in Section 2.4.2, "Unprotected DASD or Disk Failure" apply.

2.4.5 Site Loss

Site loss is the loss of a building, computer, or place to conduct business. It is usually the result of fire, flood, natural, or human-made disaster. A common misconception is that areas not affected by hurricanes, floods, or earthquakes are not at risk. Hurricanes and earthquakes account for only 8% to 10% of site loss situations. A business is more likely to shut down because of fire, terrorism, or man-made disasters, such as the evacuation of a site due to a natural gas or oil line leak ruptured by construction or a train derailment. Site loss can be limited to the computer facilities or can affect the entire business. Recovery time can range from some to never, which can and many times leads to the ultimate demise of a business.

2.5 Recovery Steps

After the failure is identified, the recovery plan is initiated. For any failure, recovery action can be outlined in general as:

1. Identify the error
2. Resolve the error
3. Recover to a known checkpoint
4. Return to the point of failure
5. Process transactions to bring the system up to a current status

These steps are identified in Figure 5 and described later in this section.




Figure 5. Steps of the Recovery Process

1. Identify the error:

The system administrator should categorize the source of the failure. Is the problem:

- As an application error?
- Hardware failure?
- IBM software defect?
- Not evident?

2. Resolve the problem:

This is a matter of having the failed item diagnosed, then repaired, replaced, or circumvented (for example, locating and replacing a piece of failed or suspect hardware, or determining a software problem and fixing it). The resolution time can involve travel time for a service representative or programmer, time to locate the problem, and time to repair the problem.

3. Recover to a known point:

This means that you need to put the system back to a known state (if required) after repairs are complete. For example, in the event of an unprotected DASD loss, this means reloading the system ASP of the AS/400 system. Restoring the complete system from two-day old backup tapes puts the AS/400 system back into a known state (which, in this case, is two days ago). In the case of a human or programming error where data is accidentally deleted or improperly updated, this can constitute reloading the system with the most recent backup of the database that was affected (which may also be two days old).

4. Return to the point of failure:

Returning the system to the point of failure happens by applying the changes and transactions that occurred since the last backup (which was restored in step 3). Applying journal changes to the database files returns the files to the point of failure. Reapply PTFs to the IBM system and product libraries if PTFs were applied since the last system save.

Note: Lack of planning can mean that transactions must be manually re-entered (if there is a paper trail from which to re-enter them). The more orphan data you have (data that has never been saved or journaled), the more impact there is on the business.

5. Process transactions to bring the system to a current status:

Even though the system may have been unavailable during the outage, business transactions may take place (for example, on a desktop PC, scratch pads, post-it notes, and so on). Depending on the business environment, these transactions may need to be entered before the system returns to full production.

Recommendation

Perform a “post mortem,” or a review of what happened.

Ask the questions:

- What caused the error?
- How could it have been prevented?

Take necessary actions to prevent a recurrence of the problem.

You can find more information on recovery planning in: *Backup and Recovery*, SC41-5304, and *So You Want to Estimate the Value of Availability*, GG22-9318.

Chapter 3. Availability Options Provided by Hardware

The AS/400 system employs a concept known as single-level storage. There is no distinction between memory and disk storage above the machine interface (MI). All storage is treated as a single, large amount of memory.

Some of the primary benefits of single level storage are:

- Increased performance (compared to other methods)
- Reduced application development time by eliminating the need to manage a disk within an application

The AS/400 system's implementation of single level storage includes automatic *scatter loading* of data across all disk units, also referred to as direct access storage device (DASD) into the *auxiliary storage pool* (ASP). Objects are spread across multiple disk arms, especially larger objects. This configuration allows multiple disk arms to service the object, potentially improving performance by servicing disk requests in parallel. It also eliminates the need to manage the disk unit where an object is located. Scatter loading makes adding new disk units easier since they are simply added to the existing pool of disk units.

This chapter describes the hardware availability options necessary to enable the protection of single level storage. Many of these options have been available since the introduction of the AS/400 system. We summarize those components for reference only and describe the enhancements from a hardware availability point of view.

The hardware options to protect *DASD availability* that are discussed in this chapter include:

- Mirrored protection
- Device parity protection
- Uninterruptible power supply

The hardware options to protect *power* and *memory* that are discussed in this chapter include:

- Uninterruptible power supply
- Battery backup
- Continuously powered main storage

Two additional topics covered in this chapter on hardware capabilities are:

- Tape device options
- Alternate installation device

3.1 Load Source Protection

Within each system unit in an AS/400 system, there are between one and four internal disk units. The first of these is designated as the load source unit. The load source unit contains the system Licensed Internal Code (LIC), dump spaces, and other information vital to the IPL and operation of the system. The load source can be protected by mirroring, device parity, or remote load source support as described in the next sections.

3.1.1 Mirrored Protection

Mirrored protection is a high availability software function that duplicates disk-related hardware components to keep the AS/400 system available if one of the components fails. It prevents a loss of data in case of a disk-related hardware failure. Mirroring is used on any model of the AS/400 system and is a part of the Licensed Internal Code (LIC).

Different levels of mirrored protection are possible, depending on what hardware is duplicated. The hardware components that can be duplicated include:

- Disk units—to provide the lowest (relative) level of availability
- Disk controllers
- Disk I/O processors (IOP)
- Buses—to provide the highest (relative) level of availability

Mirroring protection is configured by the ASP. For optimum protection, there must be an even number of components at each level of protection.

The system remains available during a disk, controller, IOP, or bus failure if the failing component and hardware components that are attached to it are duplicated.

For detailed information about mirroring functions, see *Backup and Recovery*, SC41-5304.

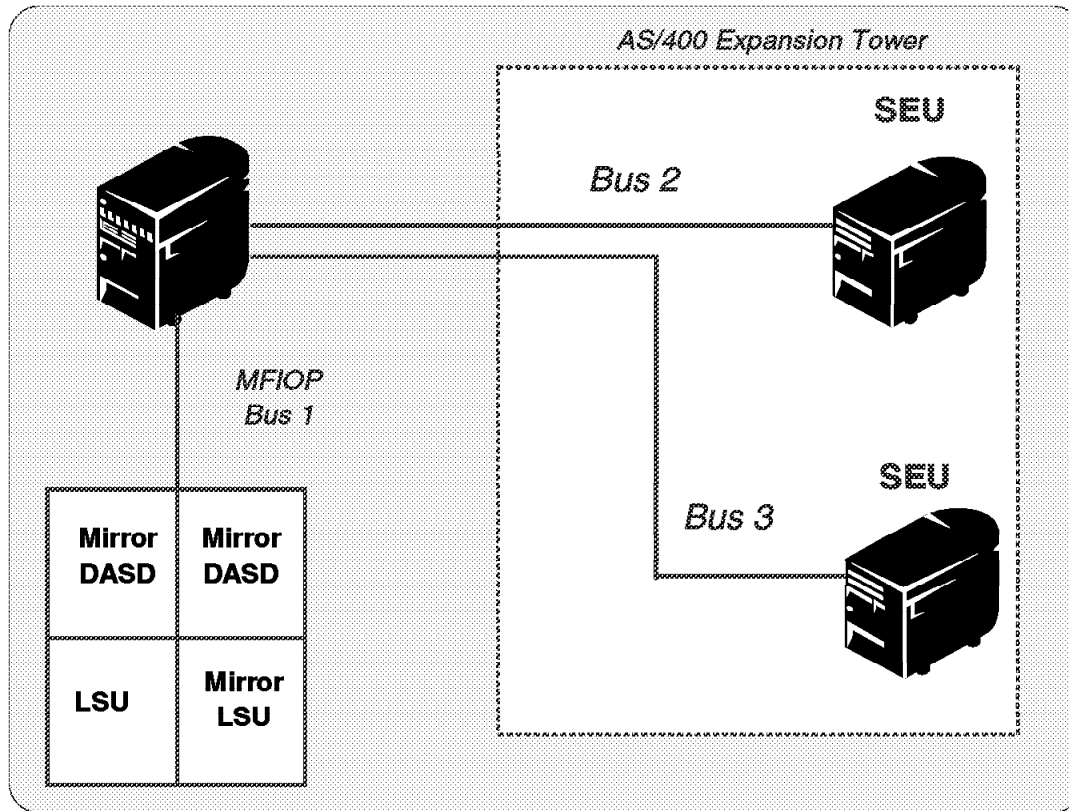
3.1.2 Standard Mirrored Protection

Standard DASD mirroring support requires that both disk units of the load source mirrored pair (unit 1) are attached to the multi-function I/O processor (MFIOIP). This option allows the system to initial program load (IPL) from either load source in the mirrored pair. The system can dump main storage to either load source if the system terminates abnormally.

However, since both load source units must be attached to the same I/O processor (IOP), controller level protection is the best mirroring protection possible for the load source mirrored pair.

3.1.3 Mirrored Load Source Protection Prior to V3R7

For systems prior to V3R7, disk-level mirroring is the standard method to protect load source devices. While this method enables the load source to be protected in the event of a disk failure or error, it does not offer the more complete protection that is available to the remainder of the system, namely IOP and bus-level mirroring. This problem is resolved with the availability of remote mirrored protection on V3R7 systems and later.



LSU = Load Source Unit

Figure 6. Standard Mirrored Load Source Protection

3.1.4 Remote Load Source Mirrored Protection on V3R7 and Later Systems

Remote load source mirroring support allows the two disk units of the load source to be on different IOPs or system buses. It provides IOP or bus level mirrored protection for the load source. The advantages of remote load source mirroring compared to standard load source mirroring are that:

- Remote load source mirroring provides IOP level or bus level mirrored protection for the load source unit.
- Remote load source mirroring allows DASD to be divided between two sites, mirroring one site to another to protect against a site disaster.

The disadvantages of load source mirroring compared to standard mirroring are that:

- A system that uses remote load source mirroring can only IPL from one DASD of the load source mirrored pair. If that DASD fails and cannot be

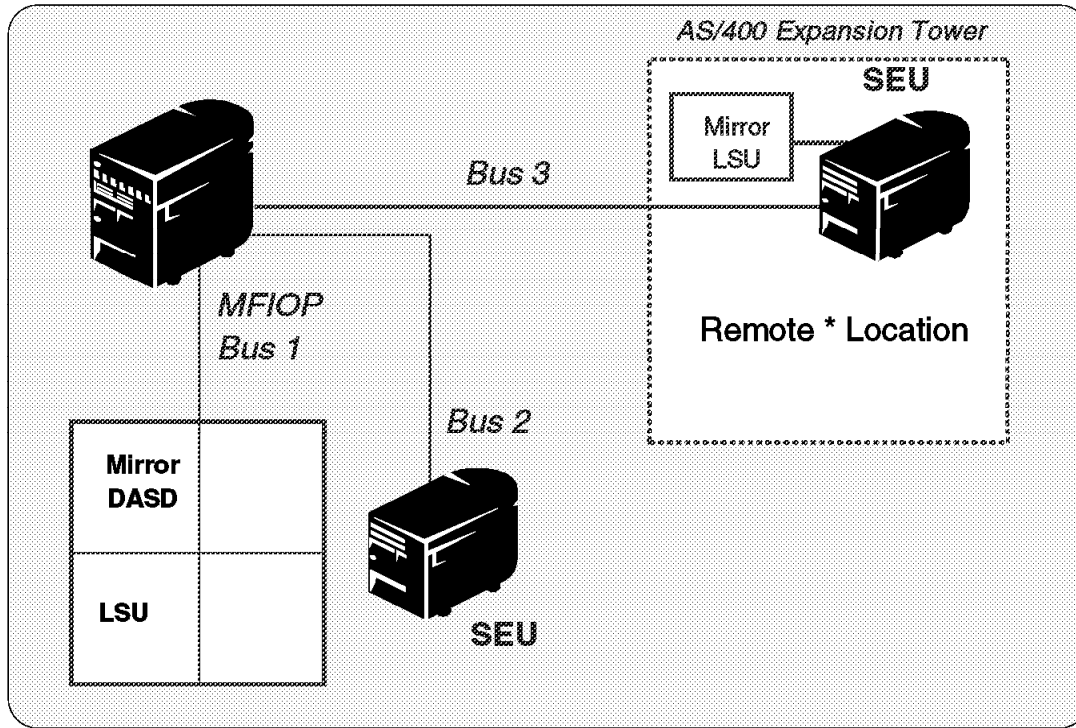
repaired concurrently, the system cannot IPL until the failed load source is fixed and the remote load source recovery procedure is performed.

- When remote load source mirroring is active on a system and the one load source that the system uses to IPL fails, the system cannot perform a main storage dump if the system abnormally terminates. This means that the system cannot use the main storage dump or continuously powered main storage (CPM) to reduce recovery time after a system crash. It also means that the main storage dump is not available to diagnose the problem that caused the system to abnormally end. The system cannot IPL or perform a main storage dump until the load source attached to the MFIOP is repaired and usable.

When combined with remote load source mirroring, the remote load source DASD function mirrors the DASD on local optical buses with DASD on optical buses terminating at a remote location. In this configuration, the entire system (including the load source) is protected from a site disaster. If the remote site is lost, the system can continue to run on the load source at the local site. Conversely, if the local DASD is lost, the remote site load source is used. If the local load source and system unit are lost, a new system unit is attached to the DASD set at the remote site. Then you can resume system processing.

Remote load source mirroring, as with standard DASD mirroring, supports mixing device parity protected disk units in the same ASP with mirrored disk units. The device parity DASD is located at either the local or the remote site. However, if a disaster occurs at the site containing the device parity DASD, all data in the ASPs that hold the device parity DASD is lost.

You can replace and synchronize the load source unit with the mirrored load source unit while the system is running.



LSU = Load Source Unit

Figure 7. Remote Load Source Protection

For some systems, standard load source mirroring remains the best choice. For others, remote load source mirroring provides important additional capabilities. Evaluate the use and needs of your system, and consider the advantages and disadvantages of each type of mirroring support to determine which is best for your availability characteristics.

For a further comparison of mirrored and device parity protection, refer to *Backup Recovery*, SC41-5304. There is also a thorough explanation of device parity, checksum protection, and mirrored protection in the redbook *AS/400 System Availability and Recovery for V2R2*, GG24-3912.

3.2 Device Parity Protection

Device parity protection is a high availability hardware function (also known as RAID-V) that protects from data loss. It allows the system to continue to operate when a disk unit fails or damage to a disk occurs. The system continues to run in an exposed mode until the damaged unit is repaired and the data is synchronized to the replaced unit.

Device parity involves calculating and saving a parity value for each bit of data. Conceptually, the parity value is computed from the data at the same location on each of the other disk units in the device parity set. When a disk failure occurs, the data on the failing unit is reconstructed using the saved parity value and the values of bits in the same location on other disks.

Logically, the implementation of device parity protection is similar to the system checksum function, except device parity is built into the hardware. System

checksum, on the other hand, is started or stopped using configuration options on the AS/400 system menu.

Note: System checksum is another disk protection method similar to device parity. Checksum is not supported on RISC systems, and we do not discuss it in this redbook. You can find information on checksum in the *Backup and Recovery*, SC41-5306, and the V2R2 redbook *AS/400 System Availability and Recovery for V2R2*, GG24-3912.

The overall goal of device parity protection is to provide high availability and protect data as inexpensively as possible. Parity protection is built into the 6502, 6512, 6532, and 6751 Input/Output Processors (IOPs). It is activated for disk units that are attached to those IOPs. It is also built into the high availability models of the 9337 Disk Array Subsystem.

For V4R2 systems, a 9751 MFIOF provides RAID support for unprotected, mirrored, or RAID protection for internal disk units as well as disk compression. The AS/400e models provide the ability to use device parity protection at bus-level distances for the load source unit. On AS/400 Models 600, 620, S10, and S20, you must order the MFIOF feature code 2726 to support device parity protection. On AS/400 Models 640, 650, S30, S40, and SB1, device parity protection is standard with MFIOF feature code 9751.

Note: The 9751 MFIOF feature code supports both device parity protection and mirroring.

Device parity protection *can*:

- Prevent your system from stopping when certain types of failures occur.
- Speed your recovery process for certain types of failures such as a site disaster or an operator or programmer error.

Device parity protection is *not*:

- A substitute for a backup and recovery strategy.
- Protection from all types of failures such as a site disaster or an operator or programmer error.

Recommendation

The system continues to run in an exposed mode until the repair operation is complete and the data is synchronized. If a failure occurs, correct the problem quickly. In the unlikely event that another disk fails in the same parity set, you can lose data.

For more information about device parity protection, see *Backup and Recovery*, SC41-5304.

3.3 Uninterruptible Power Supply

An uninterruptible power supply (UPS) provides auxiliary power to the processing unit, disk units, system console, and other devices that you choose to protect from loss of power. When you use a UPS with the AS/400 system, you can:

- Continue operations during brief power interruptions (brown outs).

- Protect the system from voltage peaks (white outs).
- Provide a normal end of operations. A normal end reduces recovery time when the system is restarted. If the system abnormally ends before completing a normal end of operations, recovery time is significant.

Normally, a UPS does not provide power to all local workstations. Nor does the UPS usually provide power to modems, bridges, or routers that support remote workstations. Consider supplying alternate power to both workgroups since the inability of worker access to information disrupts productivity. You can avoid such disruption with proper availability and recovery implementation.

Also, design your interactive applications to handle the loss of communication with a workstation. Otherwise, system resources are used in an attempt to recover devices that have lost power. Refer to Chapter 12, “Communications Error Recovery and Availability” on page 181, for more information on resources used during device recovery.

The programming language reference manuals provide examples of how to use the error feedback areas to handle workstations that are no longer communicating with the application. *Backup and Recovery*, SC41-5304, describes how to develop programs to handle an orderly shutdown of the system when the UPS takes over.

3.4 Battery Backup

Most (but not all) AS/400 models are equipped with a battery backup. Based on the system storage size, relying on a battery backup for enough time for an orderly shutdown is not sufficient.

The battery capacity typically varies between 10 and 60 minutes. The useful capacity depends on the application requirements, main storage size, and system configuration. Consider the reduction of capacity caused by the natural aging of the battery and environmental extremes of the site when selecting the battery. The battery must have the capacity to maintain the system load requirements at the end of its useful life.

Refer to *Backup and Recovery*, SC41-5304, for power down times for the Advanced Series systems. Refer to the *AS/400 Physical Planning Reference*, SA41-5109, for power down times for the AS/400 Bnn-Fnn models.

3.5 Continuously Powered Main Storage

On V3R6 systems and later, AS/400 systems are equipped with a System Power Control Network (SPCN) feature, which provides the CPM function. Upon a power fluctuation, the transition to CPM mode is 90 seconds after an initial 30 second waiting period. The internal battery backup provides sufficient power to keep the AS/400 system up for the 120 seconds until the transition to the CPM is complete. With CPM enabled, the battery provides sufficient power to shut down the system and maintain the contents of memory for up to 48 hours after a power loss without user interface or control.

The transition to CPM is irreversible. CPM interrupts the processes at the next microcode end statement and forces as many updates to disk as it can. During the next IPL, it restores main storage and attempts to complete outstanding

updates. Preserving main storage contents significantly reduces the amount of time the system requires to perform an IPL after a power loss. CPM operates outside of transaction boundaries.

You can use the CPM feature along with a UPS (or the battery backup). If the system detects that the UPS can no longer provide sufficient power to the system, the data currently in memory is put into “sleep” mode. The CPM storage feature takes control and maintains data in memory for up to 48 hours. With the CPM feature, the system automatically initiates an IPL after power is restored.

Refer to *Backup and Recovery*, SC41-5304, and the *AS/400e Series System Handbook*, GA19-5486-16, for more information on CPM requirements.

3.6 Tape Device Options

For information on what tape devices are available for each AS/400 model and the hardware and software requirements to support each model, refer to the *AS/400 Advanced Series System Handbook*, GA19-5486; *AS/400e Series System Builder*; SG24-2155, Appendix C, “Save and Restore Rates of IBM Tape Drives for Sample Workloads” on page 379; and Section 8.1, “Save and Restore Performance” on page 95, for save and restore rates.

3.7 Alternate Installation Device

On V4R1 and later systems, you can use a combination of devices that are attached on the first system bus, as well as additional buses. The alternate installation device **does not need to be attached** to the first system bus. For example, the 3590 tape drive can be positioned up to 500 meters or two kilometers away. This enables a physical security improvement, as users who are allowed access to the machine room can be different than those operating the tape drives.

You can select an alternate installation device connected through any I/O bus attached to the system. When you perform a D-mode IPL (D-IPL), you can use the tape device from another bus using the Install Licensed Internal Code display. For example, if you have a 3590 attached to a another bus (other than bus 1), you can choose to install from the alternate installation device using the Install Licensed Internal Code display and then continue to load the LIC, OS/400, and user data using the alternate installation device.

Note: Set up alternate installation device support prior to performing a D-IPL. System Licensed Internal Code (SLIC) media is necessary to perform the D-IPL that restores and installs from the tape device.

Recommendation

Before using the alternate installation device, ensure that it is defined on a bus other than system bus 1. You must enable the device. When installing from the alternate installation device, you need both your tape media and the CD-ROM media containing the Licensed Internal Code.

Some models, typically with 3590 tape devices attached, see a performance improvement when using an alternate installation device for other save and

restore or installation operations. This is caused by having the tape drive on a different IOP than the load source unit is attached to. On systems prior to V4R1, the alternate installation device is only supported using devices attached to the first system bus. The first system bus connects to the service processor IOP. Typically this is where the optical or tape devices used for installations are attached.

For step-by-step instructions to set up alternate installation device support, see *Backup and Recovery*, SC41-5304.

Chapter 4. IPL Improvements for Availability

The introduction of RISC hardware on the AS/400 system greatly improved the time to IPL. Benchmark figures show that, on average, an IPL takes 30% less time on a RISC processor than it does on a CISC processor. This improvement is due to the faster RISC hardware and other internal changes, such as an increase in page size to 4K.

Versions prior to V3R6 require an IPL to regenerate temporary addresses. Systems that created and deleted a large number of objects during normal operations required an IPL about once a week to allow addresses to be reclaimed and reused. With V3R6 and later systems, a new design and 64-bit addressing virtually eliminate the need for a regularly scheduled IPL.

For normal IPL times (where the previous system shutdown was completed to a normal end of job status), V4R1 and later systems realize up to a 50% reduction in the time it takes to do an IPL compared to V3R7 systems. Abnormal IPLs (where an abrupt shutdown causes the system to do recovery processing in either SLIC, OS/400, or both) offer up to a 30% improvement of IPL time in V4R1 and later systems compared to V3R7. Depending on the options selected prior to the IPL, an abnormal IPL can last beyond two and three hours on some systems.

In general, the performance improvements made in the IPL process span across three general phases of an IPL:

- Hardware
- System Licensed Internal Code (SLIC)
- OS/400 initialization

Within each of these IPL phases, many IPL stages (indicated by System Reference Codes (SRCs)) are substantially reduced in time. This contributes to significant improvements in the total time to perform an IPL.

This chapter helps you understand what has changed in the IPL process and how you can take advantage of the changes for improved system availability.

4.1 A Basic Understanding of an IPL

To help you understand how availability is improved by the changes made to the IPL process and changes the user can control on the AS/400 system, this section describes basic concepts regarding why an IPL is initiated, what occurs during the IPL, and the modes of an IPL.

When an IPL is performed, many cleanup and startup tasks are accomplished. The IPL process can be initiated for a variety of reasons, including:

- Starting the system after it is powered off
- Restarting the system to reset a problem
- Restarting to load a code fix (if required)
- Restarting after a crash, loop, or hang (if required)
- Restarting after a change in the hardware configuration (if required)
- Performing system maintenance such as regenerating temporary address structures (CISC only)

What happens during an IPL includes:

- Hardware diagnostics are run (as specified with IPL attributes).
- The hardware state is initialized.
- System Licensed Internal Code (SLIC) is loaded.
- The operating system is loaded and started.
- Recovery is performed as needed on user data.

The system can be started in a variety of modes:

- A = IPL the system using permanent LIC fixes only.
- B = IPL the system using temporary LIC fixes.
- C = Reserved for laboratory use.
- D = Load the microcode from removable media, not from disk.

The IPL process is shown in Figure 8, and identifies the SRCs associated with the transition to the next IPL stage.

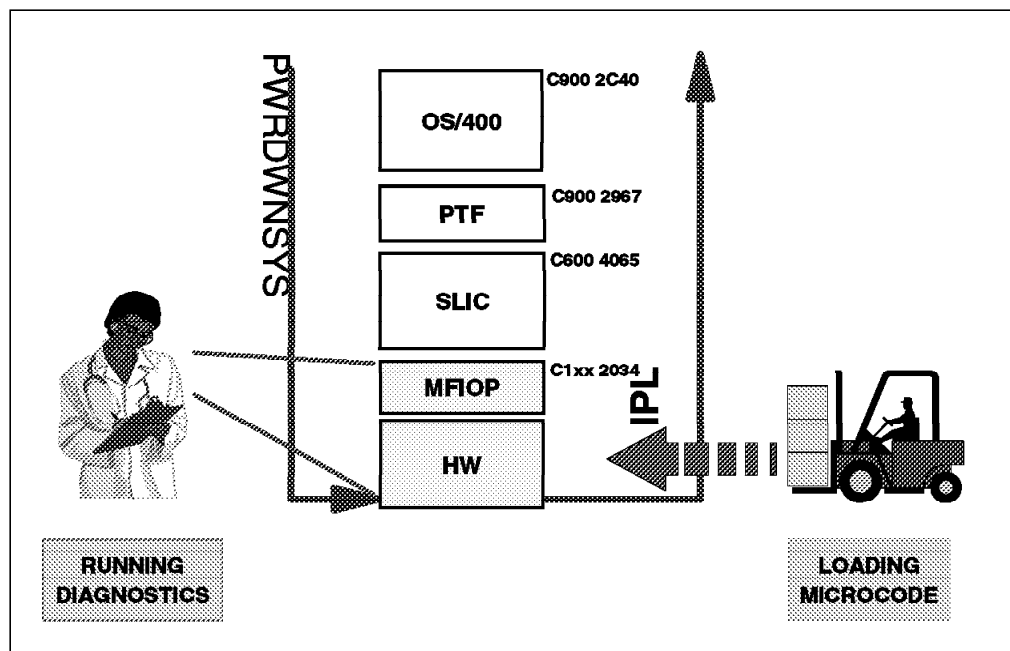


Figure 8. What is Done During an IPL

4.2 IPL Types

The IPL process can be split into various categories. You can choose to perform one of the following IPLs:

- Fast IPL with minimum hardware diagnostics
- Slow IPL where all hardware diagnostics are run
- Short programmed IPL (the default IPL)
- Full programmed IPL

Note: PTFs can be applied with a fast IPL. If the PTF is related to an MFIOP, the code is downloaded to the MFIOP first and a full IPL is done.

Specify the type of IPL you want to perform by using:

- The Change IPL Attributes (CHGIPLA) command:

When you specify HDWDIAG(*MIN) on the CHGIPLA command, the system performs a minimum, critical set of hardware diagnostics. The system runs a quick processor diagnostic check that covers approximately 90% of the hardware in approximately 10% of the time that is required for full diagnostics. The system does not perform extended main storage diagnostic or chip-to-chip circuitry tests. Most of these functions are used during the remainder of the IPL.

- The control panel on the AS/400 system unit to temporarily override the IPL attributes:

Choose function 02 (*Select IPL type/mode*) on the control panel and enter the type of IPL that you want to perform:

- F = Fast (*MIN) IPL (override current IPL attribute setting).
- S = Slow (*ALL) IPL (override current IPL attribute setting).
- V = Cancel override of current IPL attribute setting, the system value setting is used.

Selecting the IPL type using the control panel is only possible when the system is not powered on.

Find detailed instructions about using the control panel functions in *AS/400 Basic System Operation, Administration, and Problem Handling*, SC41-5206-00.

Note: The override selections stay in affect until the system is powered off at which time the override is automatically cancelled.

4.2.1.1 Fast IPL—the Default IPL

Fast IPL is the default type IPL for V4R1 and later releases of OS/400. During the hardware portion of an IPL, a minimal set of processor diagnostics is performed. A fast IPL is set using the CHGIPLA command or function 02, as discussed in the previous section.

Recommendation

Use the default value of *MIN, unless *ALL is required by service personnel.

4.2.1.2 Slow IPL

On V4R1 and later systems, you can select a slow IPL, which performs all hardware diagnostics. This is the standard in all releases prior to V4R1. A full set of IPL attribute hardware diagnostics is performed. To set a slow IPL, enter HDWDIAG *ALL on the CHGIPLA command or select function 02.

4.2.1.3 Short Programmed IPL—the Default IPL

Short programmed IPL is a term that was introduced with V4R1. It basically means that you perform a restart operation using the PWRDWNSYS RESTART(*YES *SYS) command that restarts the system LIC and the operating system. However, the system does not completely stop and restart all the hardware functions or perform an MFIOP initialization. This type of restart is the default for OS/400 unless the user enters *FULL for the restart option or specifies RESTARTTYP = *IPLA and an IPL attribute restart type = *FULL.

Usually, the IPL time for RESTART(*YES *SYS) is less than the IPL time required for RESTART(*YES *FULL). However, there is no time savings when delayed MFIOP PTFs are previously applied. Even with *SYS specified, the system

automatically does a full programmed IPL. This verifies that the new code is used during the IPL. Figure 9 on page 36 shows the difference between a RESTART type of *FULL and *SYS.

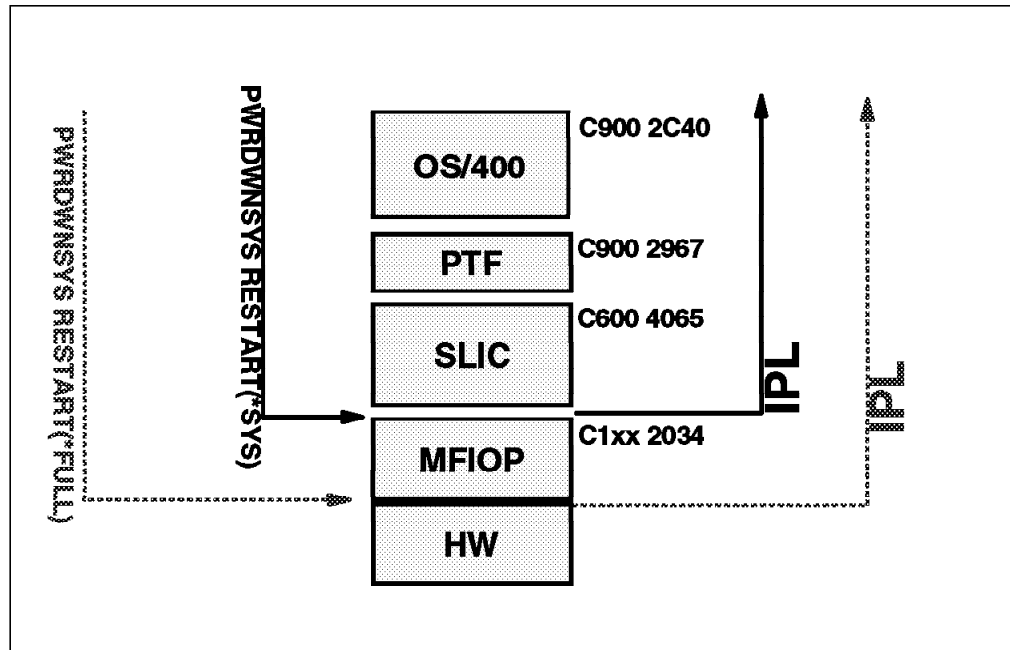


Figure 9. Restart Type *FULL versus Restart Type *SYS

4.2.1.4 Full Programmed IPL

All systems prior to V4R1 run full programmed IPL. Full programmed IPL performs a complete system restart. It indicates that the OS/400 stops all hardware functions prior to restart, and starts up all hardware functions again. MFIOP code is initialized during a full system restart.

Full programmed IPLs are used automatically by SLIC and are not selected by the operator. For example, if you have a PTF specific to the MFIOP, SLIC automatically performs a full restart. The restart option is added to the PWRDWN SYS and the CHGIPLA commands in case of unexpected problems. You select this option by specifying a restart value of *FULL or *IPLA on the PWRDWN SYS RESTART parameter.

Recommendation

Use the default value *SYS, unless *FULL is required by service personnel.

4.3 Affecting the Time to IPL

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.

There are many factors that affect IPL performance. Some of the factors that influence the duration of an IPL include:

1. Type of IPL performed:

- Normal IPL:
 - Power on IPL (cold start)
 - Programmed IPL (PWRDWNSYS RESTART(*YES))
- Abnormal IPL:
 - Abnormal system terminations cause recovery processing to be done during the IPL. The amount of processing is determined by the point in time when the system went down and how it went down, such as:
 - Function 3 to initiate a manual IPL
 - Function 8 to perform an emergency power off
 - Function 21 DST power off
 - Function 22 to retry SLIC main store dump processing
 - Function 34 to force a retry of CPM or main store dump
 - A machine check
 - A “white button” power off
 - A power outage with continuous power main storage (CPM)
 - A power outage with no protection.

Note: Be aware that the use of ENDJOBABN on a job causes the subsequent IPL to be an abnormal IPL, even if OS/400 terminates normally.

2. Hardware configuration:

- CPU—model and feature code
- DASD—amount, type, number of arms on the MFIO, protection used, size of objects stored, and number of ASPs
- Amount and type of IOPs
- Amount of main storage
- Number of device descriptions on the system
- Number of towers

3. Software configuration:

- Total number of jobs on the system
- Number of spooled files
- Number of user profiles
- Number of libraries
- Use of journal, commit, and system managed access path protection (SMAPP)
- Database size and characteristics
- Type of active workload

IPL duration highly depends on the mode and type of IPL performed and the hardware and software configuration. However, there are tasks that reduce the amount of time required for the system to perform an IPL. As a user, you can avoid a lengthy IPL operation by remembering to:

- Reduce the amount of rebuild time for access paths during an IPL by using SMAPP. *Backup and Recovery*, SC41-5304, describes this method for protecting access paths from long recovery times during an IPL, as does Section 4.4, “System Managed Access Path Protection” on page 38.
- Control the level of hardware diagnostics that are run during an IPL. When you specify the default (HDWDIAG(*MIN)) on the CHGIPLA command, the system performs 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 introduced to the system.

- Perform miscellaneous cleanup of access paths. You can use the Edit Rebuild of Access Paths (EDTRBDAP) display to postpone the rebuild of selected access paths until after the IPL.
- Keep the number of journals to a minimum and control the number of journal changes that request a delete of the receivers (DLTRCV(*YES)). A delete of receivers is done during the IPL.
- Keep the number of CHGJRN operations performed for each journal with manage receivers (MNGRCH(*SYSTEM)) to a minimum.
- Reduce the number of device descriptions. Remove any obsolete device descriptions from the system.
- Defer unnecessary startup processing for devices, subsystems, and jobs from the IPL startup job stream.

Note: Using APPN instead of APPC causes a vary on at first use. This avoids using IPL time to vary on the description.

- Minimize the number of jobs and job structures on the system. The best way to do this is to remove unnecessary spooled files. To find the current number of job structures needed on the system, use the Display Job Tables (DSPJOBTL) command described in Section 10.3, “Display Job Tables” on page 143. Use the CHGIPLA command described in Section 4.5, “Changing IPL Attributes” on page 42 to affect the duration of IPL.
- Reduce the number of systems network architecture distribution services (SNADS)-specific objects. When users initiate mail requests at a faster rate than the mail is sent or delivered, an unbalanced condition occurs. SNADS takes more time than normal to clean up. This cleanup is completed only during an abnormal IPL or RCLSTG.
- Avoid starting the IPL with function 03 (to initiate a manual mode IPL) or function 08 (to initiate an emergency power off) since they may cause a long lasting IPL.

4.4 System Managed Access Path Protection

Recovery time for rebuilding access paths can be long-running and cause extended outages after an abnormal system end. One way to reduce the recovery time is to use the System Managed Access Path Protection (SMAPP) feature. SMAPP is based on a user selected target recovery time.

SMAPP offers system monitoring of potential access path rebuild time. You select a time threshold for rebuilding. This is the maximum amount of time in minutes that is allowed for rebuilding keyed access paths during an IPL recovery. Recovery time can be specified for each ASP. Or use one number for the entire system. SMAPP starts and stops the journaling of system selected access paths automatically and dynamically to meet the target time threshold for access path recovery time. SMAPP is complementary to any existing database journaling activity and, where appropriate, shares existing journal receivers.

All keyed access paths on the system are scanned. SMAPP calculates the total time required to rebuild these access paths and compares the calculated time to the target recovery time. Based on the difference between the two times, SMAPP selects which keyed access paths to protect. SMAPP journals the

selected access paths under-the-covers to ensure that the target recovery times are not exceeded in the event of an abnormal system termination. The actual time required for rebuilding keyed access paths may vary slightly from the target recovery time.

Use the Display Recovery for Access Paths (DSPRCYAP) or Edit Recovery for Access Paths (EDTRCYAP) commands to view and specify the target recovery time.

Note: The target recovery time field is blank if an internal system failure keeps the system from determining the access path recovery time.

The shorter the target recovery time, the greater the number of access paths that must be protected and the greater need for additional system resources to perform internal journaling.

The default system-wide access path recovery time for SMAPP is 150 minutes. This means that SMAPP protects the system so that there is generally no more than 150 minutes of access path rebuild time during an IPL after an abnormal termination. SMAPP assumes the responsibility of providing the necessary amount of access path protection. No user intervention is required because SMAPP manages the entire journal environment.

SMAPP coexists with current journal functions and uses any suitable existing receiver. For access paths journaled under SMAPP support, the access path journal entries are placed into internal journal receivers unless the associated file is also journaled. In that case, SMAPP journal entries are placed into the associated journal receiver. Where no suitable receiver exists, SMAPP builds its own internal receiver. OS/400 users cannot access these SMAPP entries.

On user-defined journal receivers, a SMAPP entry appears as a missing entry sequence number.

Note: The SMAPP function runs below the operating system and requires very little overhead. Along with design enhancements for V4R2, there are PTFs for all OS/400 releases back to V3R1 that enhance the performance from V3R1 and later versions. Contact IBM AS/400 Software Support to find out what PTFs are needed for the release of your system.

4.4.1 SMAPP Tasks

When SMAPP is enabled, the WRKSYSACT command shows three SMAPP tasks. They appear as system jobs (they do not appear under a subsystem). The tasks are:

- #JOTUNT

This task performs background tuning performance to obtain the target recovery time.

- #JOEVAT

This task evaluates and calculates the rebuild time for an access path.

- #JOIJSS

This task implicitly performs start and stop protection for physical files and access paths.

4.4.2 Performance Considerations when SMAPP is Activated

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 minimizes the performance impact, while at the same time provides a reasonable and predictable recovery time and protection for key access paths. For many environments, even 60 minutes of IPL recovery time has negligible overhead. Although SMAPP may start journaling access paths, the underlying SMAPP support is much less resource intensive in terms of performance than explicit journaling support.

Note that as the target access path recovery time is lowered, the performance impact from SMAPP increases. Balance your recovery time requirements against the system resources required by SMAPP.

In general, SMAPP adds no more than 3% to CPU usage. If you experience higher values, investigate the three SMAPP tasks listed in Section 4.4.1, "SMAPP Tasks" on page 39.

To improve disk activity along with SMAPP, we recommend the following steps:

- Since SMAPP minimizes disk writes, change the attributes of any physical or logical files that use force write ratio. Do this with the Change Physical File (CHGPF) command:

```
CHGPF FILE(library-name/file-name) FRCRATIO(*NONE)
```

- Do the same for the associated logical file. Enter the command:

```
CHGLF FILE(library-name/file-name) FRCRATIO(*NONE)
```

This command allows the system to determine when to write records to disk and eliminate unnecessary disk writes.

- Reduce SMAPP overhead by installing disk units that support write cache, since SMAPP writes to the 10 fastest disk arms in an ASP.

For more information on SMAPP and explicit access path journaling, see *Backup and Recovery*, SC41-5304.

4.4.3 Modifying SMAPP

Although the default level of SMAPP protection is sufficient for most customers, some customers need a different level of protection. We recommend that you modify the SMAPP settings. The important variables are:

1. The number of key changes
2. The number of unprotected access paths

For those users who have experienced abnormal IPL access path recovery longer than 150 minutes, we advise that you experiment by varying the amount of protection. Too much protection causes undesirable delays in the IPL.

Recommendation

During the busiest part of the day, set the recovery value temporarily to *NONE, display the estimated access path recovery time with the DSPRCYAP command, and decide whether this meets your requirements. If not, change the recovery time to whatever it needs to be with the CHGRCYAP command.

By understanding your system requirements and experimenting to find what value meets these requirements, you can decide on an optimum SMAPP setting.

The following displays show examples of the CHGRCYAP, EDTRCYAP, and DSPRCYAP commands.

```

Chg Recovery for Access Paths (CHGRCYAP)

Type choices, press Enter.

System recovery time . . . . . 150      10-1440, *SAME, *SYSDFT...
ASP recovery time:
Auxiliary storage pool ID . .  1        1-16
Recovery time . . . . . *NONE      10-1440, *SAME, *NONE, *MIN

Auxiliary storage pool ID . .  3        1-16
Recovery time . . . . . *NONE      10-1440, *SAME, *NONE, *MIN

Auxiliary storage pool ID . .  12       1-16
Recovery time . . . . . *NONE      10-1440, *SAME, *NONE, *MIN
+ for more values

```

Figure 10. The Change Recovery Access Paths Command

```

Edit Recovery for Access Paths                                SYSTEMXX
                                                            02/27/98 20:27:26
Estimated system access path recovery time . . . :    35 Minutes
Total disk storage used . . . . . :    .081 MB
% of disk storage used . . . . . :    .000

Type changes, press Enter.
System access path recovery time: . . .      150 *SYSDFT, *NONE, *MIN
                                                *OFF, Recovery time

-----Access Path Recovery Time----- --Disk Storage Used---
ASP   Target (Minutes)   Estimated (Minutes)   Megabytes   ASP%
  1       *NONE           12                   .098       .000
  3       *NONE           0                    .032       .000
 12       *NONE           0                    .024       .000

```

Figure 11. The Edit Recovery for Access Paths Command

Note: The system throttles the % of disk storage used to a minimal value. The .000 value represented in Figure 11 does not represent an actual measurement.

```

Display Recovery for Access Paths          SYSTEMXX
                                           12/10/97 20:27:22
Estimated system access path recovery time . . . :    35 Minutes
Total disk storage used . . . . . :    .081 MB
% of disk storage used . . . . . :    .000
System access path recovery time . . . . . :    150

-----Access Path Recovery Time----- --Disk Storage Use
ASP   Target (Minutes)   Estimated (Minutes)  Megabytes   ASP
  1      *NONE             12                .098       .000
  3      *NONE             0                 .032       .000
 12      *NONE             0                 .024       .000

```

Figure 12. The Display Recovery Access Paths Command

4.5 Changing IPL Attributes

Use the Change IPL Attributes (CHGIPLA) command to change the settings of the attributes used during an IPL. With this command, you can affect the time it takes to IPL the system by altering the following values:

- Restart options
- Keylock position
- Whether the job tables are compressed
- Whether the product directory is rebuilt
- Whether to clear out the output queues
- Whether to restart the printers

The restart type, hardware diagnostics, and job table verification options are highlighted in Figure 13 on page 43, and discussed in the following sections.

Note: These changes take affect for the *subsequent IPL* only.

```

Change IPL Attributes (CHGIPLA)

Type choices, press Enter.

Restart type . . . . . *SYS      *SAME, *SYS, *FULL
Keylock position . . . . . *SAME   *SAME, *NORMAL, *AUTO...
Hardware diagnostics . . . . . *MIN    *SAME, *MIN, *ALL
Compress job tables: . . . . . *NONE   *SAME, *NONE, *NORMAL...
Check job tables . . . . . *ABNORMAL *SAME, *ABNORMAL, *ALL, *SYNC
Rebuild product directory . . . *NONE   *SAME, *NONE, *NORMAL...
Mail Server Framework recovery *NONE   *SAME, *NONE, *ABNORMAL
Clear job queues . . . . . *NO     *SAME, *YES, *NO
Clear output queues . . . . . *NO     *SAME, *YES, *NO
Clear incomplete joblogs . . . *NO     *SAME, *YES, *NO
Start print writers . . . . . *YES    *SAME, *YES, *NO
Start to restricted state . . . *NO     *SAME, *YES, *NO

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 13. The Change IPL Attributes Command

4.5.1 Restart Type (RESTART) Parameter

The restart type (RESTART) parameter specifies the point from which the IPL restarts when RESTART(*YES) or RESTART(*YES *IPLA) is entered on the Power Down System (PWRDWN SYS) command. Entering *SYS rather than *FULL reduces the time required to restart the system. The initial (shipped) value for this parameter is *SYS. The possible values are:

- *SYS** The operating system restarts. The hardware restarts only if a PTF requiring a hardware restart is applied. The same applies to LIC PTFs.
- *FULL** All portions of the system, including the hardware, restart.

Note: With the *SYS option, the MFIOP is not reset. Therefore, communication error or hardware recovery does not end until a *FULL IPL is completed. Also, certain status conditions in the hardware IOPs are not reset without a full IPL.

Recommendation

Use the default value *SYS unless *FULL is required by service personnel.

4.5.2 Hardware Diagnostics (HDWDIAG) Parameter

The hardware diagnostics (HDWDIAG) parameter applies when the system performs an IPL from a powered off state. It indicates whether to perform certain hardware diagnostics during the IPL. The user cannot modify the list of these diagnostics. The initial (shipped) value for this attribute is *MIN. The possible values are:

- *MIN** The minimum set of hardware diagnostics runs.
- *ALL** All hardware diagnostics run.

Recommendation

Choose the *ALL option at least annually, since certain service actions require an *ALL function to be performed. Normally use the default value *MIN, unless service personnel require the *ALL function.

For additional information on how to improve IPL performance, refer to *AS/400 Basic System Operation, Administration, and Problem Handling*, SC41-5206-01 Refer to Chapter 12, “Communications Error Recovery and Availability” on page 181 for an additional discussion on communications configuration objects and an IPL.

4.5.3 Check Job Tables

Many of the improvements in the time to IPL have been made in the area of verifying the job tables for damage. The check job tables function is run during abnormal IPLs to determine when to perform certain damage checks on the job tables. By default, the IPL process checks the job tables if the IPL is abnormal. You can also specify *SYNC to synchronously check for damage during all IPLs.

Recommendation

For larger users or those with volatile job activity, specify *SYNC on the CHKJOBTL command to enable job-table checking during all IPLs.

Refer to the *CL Reference*, SC41-5722, for more information on the CHKJOBTL parameter and CHGIPLA command.

4.6 Marking the Progress of an IPL with SRC Codes

SRC codes indicate which function is performing by the processor at various points of an IPL. Monitor the codes displayed on the control panel during an IPL to get an indication of what is being done.

There is no estimated time that each step should take, since each system is unique and the timings differ on each IPL on any given system. However, when you build a history of the times each IPL step takes on your system, you build a history of what is “normal” for your system. Subsequently, you can influence the time to IPL by some of the actions described in Section 4.3, “Affecting the Time to IPL” on page 36.

Note: Not all SRC codes are listed. A suggested action is listed below the description for some of the listed SRCs.

Codes displayed in the pre-OS/400 IPL stage include:

- C6004272—ASP overflow recovery
- C6004025—Authority recovery
- C6004026—Journal recovery
- C6004027—Database recovery

Codes displayed in the OS/400 IPL stage include:

- C9002810—Reclaim machine context
- C9002825—Convert the Work Control Block Table (WCBT)
 - Reduce the size of the WCBT

- Perform a scratch installation to eliminate this step
- C9002830—Validate system values, check for duplicate device descriptions
 - Eliminate unnecessary device descriptions
- C90028C0—Prepare job instructions
- C9002910—Start system logging of messages
 - Reduce the number of messages in QHST and QSYSOPR
- C9002920—Library cleanup
 - Control the number of libraries
 - Eliminate temporary job structures
- C9002925—POSIX directory cleanup
- C9002930—Set up database cross reference
- C9002940—Set up and vary on QCTL and QCONSOLE
- C9002960—IPL sign on and PTF processing
- C9002965—Software management services initialization
 - Control the number of libraries and licensed program products on the system
- C9002970—Database, journal, and commit functions
 - Control the number of journals that specify MNGRC(*SYSTEM)
 - Avoid using system request to cancel out of jobs when they are in database operations
- C9002990—IPL performance tuning of machine and base pools
 - Eliminate unnecessary device, controller, line and network interface descriptions
- C90029A0—Prepare the system control block structure
 - Control the number of libraries
 - Manage the QUPSMGQ and QINACTMSGQ to clear unneeded messages
 - Reduce the size of the WCBT
- C90029B0—Spool initialization
 - Delete unnecessary spool files
- C90029C0—Work Control Block Table initialization
 - Reduce the size of the WCBT
- C9002A80—Miscellaneous cleanup
- C9002A85—Create shared activation group for POSIX
- C9002A90 – Start system jobs
- C9002A95—Cleanup of Work Control Block Table during an abnormal IPL
 - Reduce the size of the WCBT
- C9002AA0—Cleanup for database, journaling, and commitment control
 - Manage the number of journals created by specifying DLTRCV(*YES)
 - Control the number of files, members, and access paths
 - Manage the use of commitment control and two phase commit operations
- C9002AA5—Recover POSIX directories
- C9002AB0—Miscellaneous cleanup including access paths
 - Manage the number of files created with *IMMED or *DLY access path maintenance that specify a recovery value of *IPL
 - Start access path journals with the STRJRNAP command
 - Balance the value specified for EDTRCYAP to provide an acceptable IPL and run time performance
- C9002B20—Process the Resource Configuration Record (RCR)
- C9002B30—Initialize and start QCLUS
 - Eliminate unnecessary device descriptions
 - Use APPN rather than APPC
- C9002B40—Device configuration
 - Eliminate unnecessary device descriptions

- Limit the number of device descriptions which specify *YES for online at IPL
- C9002C20—QSNADS object recovery
 - Reduce the number of SNADS objects (mail)
- C9002C30—Create temporary job structures
- C9002C40—Work Control Block Table cleanup for a normal IPL
 - Reduce the size of the WCBT
- C9002C50—Recovery of office objects
 - Reduce the number of documents and folders (DLOs)

Note: A suggested improvement for several of the IPL steps involves reducing the size of the WCBT (the internal system object known as the Work Control Block Table). The size of the WCBT can be kept to a minimum when you:

1. Clean up jobs that have ended and in which spooled output is no longer needed. Jobs which have completed are removed from the WCBT when they no longer have spooled output associated with them.
2. Limit the number of available (unused) entries in the WCBT. The number of available entries is affected by the system values QTOTJOB and QADLTOTJ.

Refer to Section 10.2.1, “Work Control Block Table Cleanup” on page 143, for more information on WCBT cleanup.

Note: Use the report produced by the QWCCTREC tool as a reference to indicate when these SRCs appear. Refer to Appendix B, “Evaluating the Time to IPL” on page 371 for a sample report.

You can find additional information on the stages of an IPL and SRC codes in the *AS/400 Licensed Internal Code Diagnostic Aids - Volume 1, LY44-5900*, reference manual.

4.7 IPL Benchmarks

To understand the improvements made in the time to IPL, a summary of the benchmark used in IBM testing is provided.

Note: The information that follows is based on performance measurements and analysis done in the IBM AS/400 laboratory. Actual performance on any given system varies significantly from what is described here (see Section 4.3, “Affecting the Time to IPL” on page 36).

Normal IPLs were measured from power on through the console signon display.




Figure 14. Normal IPL Time Improvement

Note: Remember that some relatively long running tasks, such as starting TCP/IP or varying on an Integrated PC Server, take place after the IPL is completed. Seeing the signon display on the console does not mean that all other workstations are up and ready for users to sign on.

For abnormal IPLs, the benchmark consists of bringing up a database workload and letting it run for a specified period of time. Once the workload stabilizes, select a function 22 to force a main store dump. This environment simulates a loop or hang situation. The dump is copied to DASD with the Auto Copy function available in V4R1 through System Service Tools. Once the dump is copied, the system completes the remaining IPL without user intervention. This is also possible by using the Auto Copy function and switching the key to normal mode shortly after you select function 22.

Benchmark time is measured from the time you select function 22 to the time the console signon display appears.

4.7.1 Benchmark Configuration

This section describes the standards for the benchmark.

1. The hardware and software environment included:
 - 530-2162 (4-way) with 4GB main storage and 164GB DASD
 - Ten towers
 - Four thousand local workstations
 - Two token-ring cards
2. The system had 20000 output queue jobs, 10000 jobs on job queues, 3450 user profiles and 2700 user libraries. The application and database environment that were used consisted of:

- One library with 200 physical files and 20 logical files
- One library with 100 physical files and 10 logical files
- One library with 50 physical files and 10 logical files
- One library with 100 physical files and 20 logical files
- One library with 20 physical files and 5 logical files
- One library with 150 physical files and 10 logical files
- One library with 100 physical files and 100 logical files
- One library with 15 physical files and 1 logical files
- One library with 250 physical files and 10 logical files
- One library with 15 physical files and 5 logical files
- One library with 10 000 physical files and 200 logical files
- CPW database
- Three ASPs—Journal Receivers are in ASP 2 and ASP 3

All physical files were explicitly journaled. The logical files were protected by setting SMAPP to *MIN. The journal receivers resided in user ASP 2 and ASP 3.

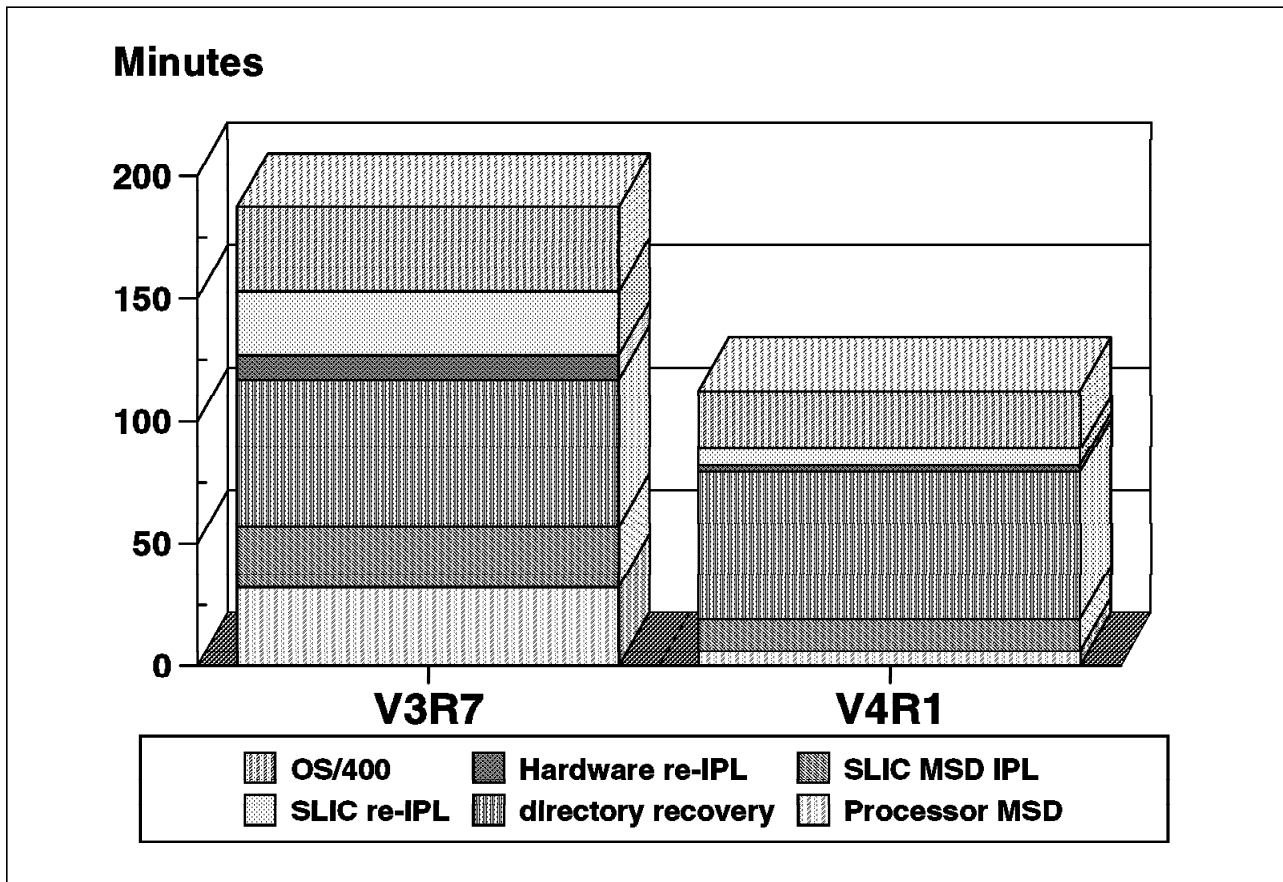


Figure 15. Abnormal IPL Time Improvement

Note: The benchmark results are based on a controlled lab environment and are for your reference only. Depending on your system configuration and software setup, your results may vary significantly.

Chapter 5. Save and Restore for Availability and Recovery

The fundamentals for a backup and recovery plan include instructions on saving information in a manner that allows for recovery if planned and unplanned outages occur. The ultimate objective is to provide a reliable backup and recovery plan that requires a minimum save and restore window size.

This chapter discusses save and restore options to enable a more effective backup and recovery plan.

5.1 SAVxxx and RSTxxx and Flexibility

System management and operations are improved for V4R2 and later systems to allow more flexibility. You can customize backup plans further by saving objects in different combinations to avoid saving those objects that do not require frequent backups. For example, the SAVLIB, SAVOBJ, SAVCHGOBJ commands, and the QSRSAVO API allow generic values for the LIB parameter. To back up all user libraries beginning with the characters DMT, enter:

```
SAVLIB LIB(DMT*)
```

SAVLIB, SAVOBJ, SAVCHGOBJ, and the QSRSAVO API also support the OMITLIB parameter for omitting libraries that do not require a save. For example, to back up all user libraries except the library named DAN, enter:

```
SAVLIB LIB(ALLUSR*) OMITLIB(DAN)
```

Up to 300 specific or generic named objects can be omitted from a save. Since less information is saved, time to save decreases, which improves availability while not sacrificing ease-of-use.

Note: On systems prior to V4R1, the OMITLIB parameter is supported only if *ALLUSR, *IBM, or *NONSYS is specified.

Two additional options providing flexibility on save functions are:

- Sequence numbers up to 16 777 215.

With increased offline media capacity, more objects are saved. The larger sequence numbers on save related commands support the growth in tape technology and the growing customer environment by enabling sequence number values up to 16 777 215. The previous limit was 9 999.

Note: For V4R1 and earlier systems, the largest sequence number allowed is 65 535 for the following commands:

- CPYTOTAP
- CPYFRMTAP
- CRTTAPF
- CHGTAPF
- OVRTAPF
- CHKTAP
- DMPTAP
- DUPTAP
- DSPTAP

On V4R2 systems, the previously listed V4R1 commands, and the following, allow sequence numbers up to 16 777 215:

- SAV/RST
- SAVSECDTA
- RSTUSRPRF
- CPYPTF
- LODPTF
- LODRUN
- RTVOBJD
- SAVSAVFDTA
- SAVCFG/RSTCFG
- SAVCHGOBJ
- SAVDLO
- SAVLIB/RSTLIB
- SAVLICPGM/RSTLICPGM
- SAVOBJ/RSTOBJ

Note: The maximum sequence number allowed on Document Library Services commands (SAVDLO and RSTDLO) remains at 65535. This is due to the architecture referenced by user applications. Sequence numbers specified greater than 65535 are handled by subtracting the remainder from 64K. For example, 65536 has an off-line sequence number of 00.

- The *ERR value is supported for the information type (INFTYPE) parameter. This allows for faster save processing. In addition, the outfile used in the output operation contains a record for each object that fails to save and a message indicating why the save failed. On systems prior to V4R2, only successfully saved object information is recorded in the outfile.

5.1.1 Concurrent Save with Generic OMITLIB Example

The save commands allowing the user to specify generic library names make it easier to run concurrent SAVLIB LIB(*ALLUSR) requests to multiple tape drives. By specifying generic omit values, multiple tape drives can be used concurrently. Each tape drive can contain different portions of the backup. For example, libraries can be split into two saves based on how they are named.

The following example shows a concurrent SAVLIB request. The two SAVLIB commands save all libraries that begin with the letters A through Z on to two separate tape drives.

```
SAVLIB LIB(*ALLUSR) DEV(TAP04) OMITLIB(A* B* C* ...L*)
SAVLIB LIB(*ALLUSR) DEV(TAP05) OMITLIB(M* N* O* ...Z*)
```

After these SAVLIB commands are run, TAP04 tape drive contains all user libraries *not* starting with A through L, and TAP05 contains all user libraries *not* starting with M through Z. Together the set of tapes contain all libraries beginning with A through Z.

Note: If you implement an example based on the beginning letter of an object's name, remember to consider objects that start with special characters, such as &, #, and so on.

5.2 Save While Active and Object Locks

The Save While Active (SWA) function helps improve availability by allowing save operations without ending jobs or subsystems. Objects are saved even while they are in use. SWA reduces the amount of time applications are unavailable to users due to backup requirements.

The SWA function offers these benefits:

- Objects with logical dependencies in the same library reach the same checkpoint.
- Objects in the same library that are journaled to the same journal reach the same checkpoint.
- Objects in the same library being processed under commitment control reach the checkpoint at a commitment boundary.

These benefits make recovery easier to manage due to the number of checkpoints involved.

Beginning with V4R1, many of the restrictions for SWA are removed. Most importantly, the following results occur:

- Once a save checkpoint is reached, locks for most object types are removed or reduced. On systems prior to V4R1, post-checkpoint processing retains *SHRUPD, *SHRRD, or *SHRNUP locks (depending upon the object type).

With a *SHRRD or no lock option, more applications can take advantage of the post processing save window.

- More applications can restart after the checkpoint CPI3712 message indicating that SWA checkpoint processing is complete.
- After a checkpoint is reached, you can:
 - Remove or rename database members
 - Delete, move or rename most objects, including files
 - Clear physical file members
 - Start a subsystem
 - Perform selected S/36 environment operations, such as LIBRLIBR

Types of SWA checkpoint processing are:

1. Synchronized libraries
2. Library
3. Command
4. System-defined

Figure 16 on page 52 provides an overview of the save-while-active process. The “Tn” labels marking the steps in the SWA process are outlined following the figure.

Save-While-Active

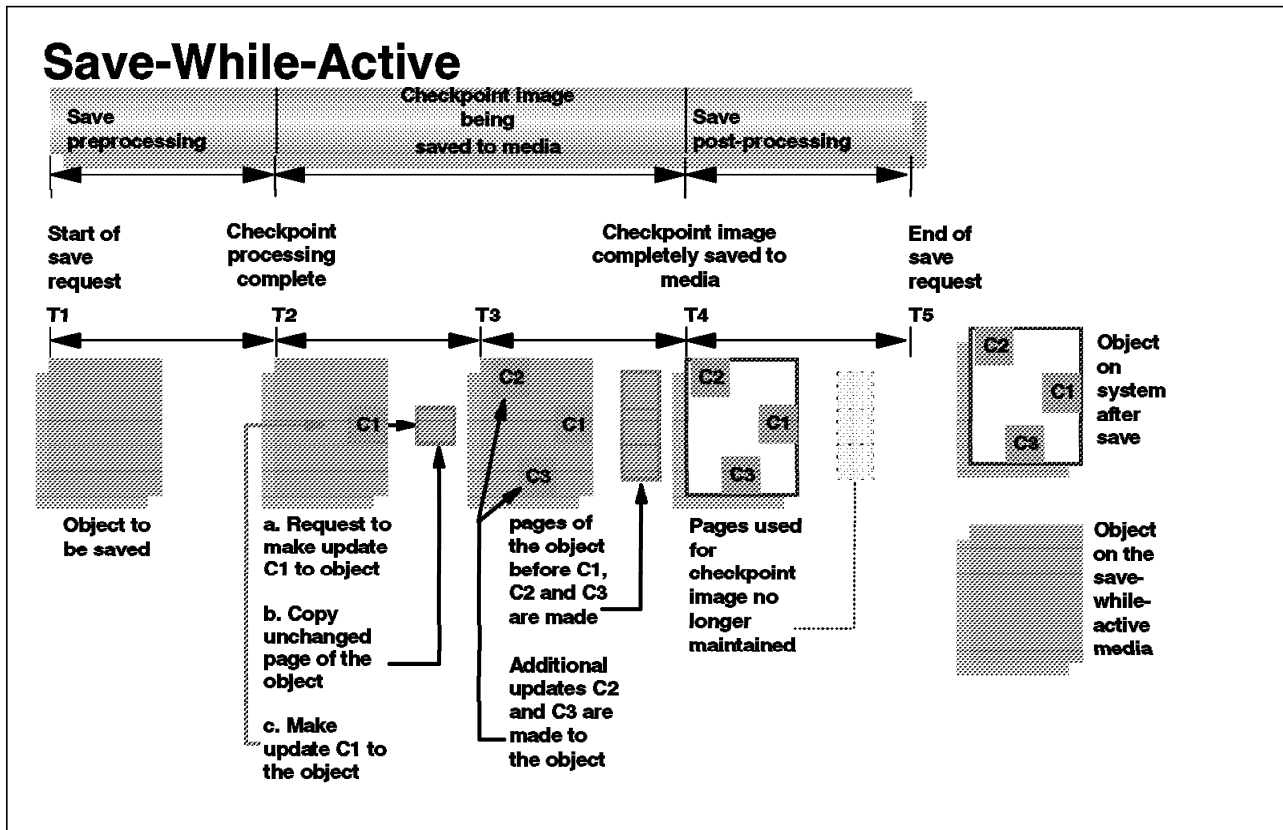


Figure 16. An Overview of the Save-While-Active Process

1. Freeze the object and establish a checkpoint as shown in Figure 16. Time T1 is the save preprocessing phase of the save-while-active operation. At the end of time T1, the object has reached a checkpoint.
2. Object changes activate a “shadow” or “side” file. Time T2 shows an update to the object, referred to as C1, while the object is saved to the media.
 - a. A request is made to update C1.
 - b. A copy of the original pages are sent to the “side” file first.
 - c. The change is made to the object. The original page copied is part of the checkpoint image for the object.
3. The save process reads from the “side” file. Time T3 shows two additional changes, C2 and C3, made to the object. Additional change requests that are made to the pages of the object already changed for C1, C2, or C3 do not require any additional processing. The requests are made with respect to the checkpoint image of the object. At the end of time T3, the object is completely saved to the media.
4. Time T4 shows that the copied pages for the checkpoint image of the object are no longer maintained because they are no longer needed.
5. Time T5 shows that the object on the system has the C1, C2, and C3 changes. The copy or image of the object saved to the media does not contain those changes.

Note: No CLRPFM, object deletions, renaming, moves, or subsystem starts may occur during this process. On V4R1 or later, these can occur after checkpoint processing is complete.

5.2.1 Save While Active Considerations

Despite the changes for SWA processing that enhance the save process, be aware of the following considerations that remain for managing SWA operations:

The “library in use” message is issued when:

- Additional save or restore operations are performed on objects or libraries that are being saved.
- A delete, rename, or reclaim operation is attempted on a library that is being saved.
- PTFs are loaded, applied, removed, or installed on a library on which objects are being saved.
- Licensed programs are saved, restored, installed, or deleted from objects that are being saved.

The “object in use” message is issued when:

- A CHGPF is attempted when SRCFILE, ACCPTHSIZ, NODGRP, PTNKEY, or an alter statement is entered for SQL operations.
- A delete or move operation is attempted on a journal receiver.
- Journal recovery is attempted involving any receiver.
- A detach or attach operation is attempted on a journal.
- A delete or move operation is attempted on a journal with an attached receiver.
- A delete, rename, or move operation is attempted on a PRDLOD object.

As always, if media errors occur on the tape during a SWA operation, there is no recovery. Check for successful job completion and start the SWA save again when media errors occur.

For more information about save-while-active, see *Backup and Recovery*, SC41-5304.

5.2.2 Save While Active and Target Release

On systems prior to V4R2, the Save While Active (SAVACT) option is permitted when the target release (TGTRLS) parameter contains the value *CURRENT. On V4R2 or later systems, target release support allows *CURRENT, V4R2M0, V4R1M0, or *PRV, V3R7M0, and V3R2M0 as valid values for the save commands.

See Chapter 6, “Save and Restore Considerations for Mixed Release Environments” on page 75, for more information on target release.

5.3 Omitting Objects on a SAVSYS Operation

On V4R1 and later, an OMIT parameter is available on the SAVSYS command. This parameter allows the user to omit:

- *CFG—Configuration objects
- *SECDTA—All security objects (user profiles and security information)

These options are highlighted in Figure 17 on page 54.

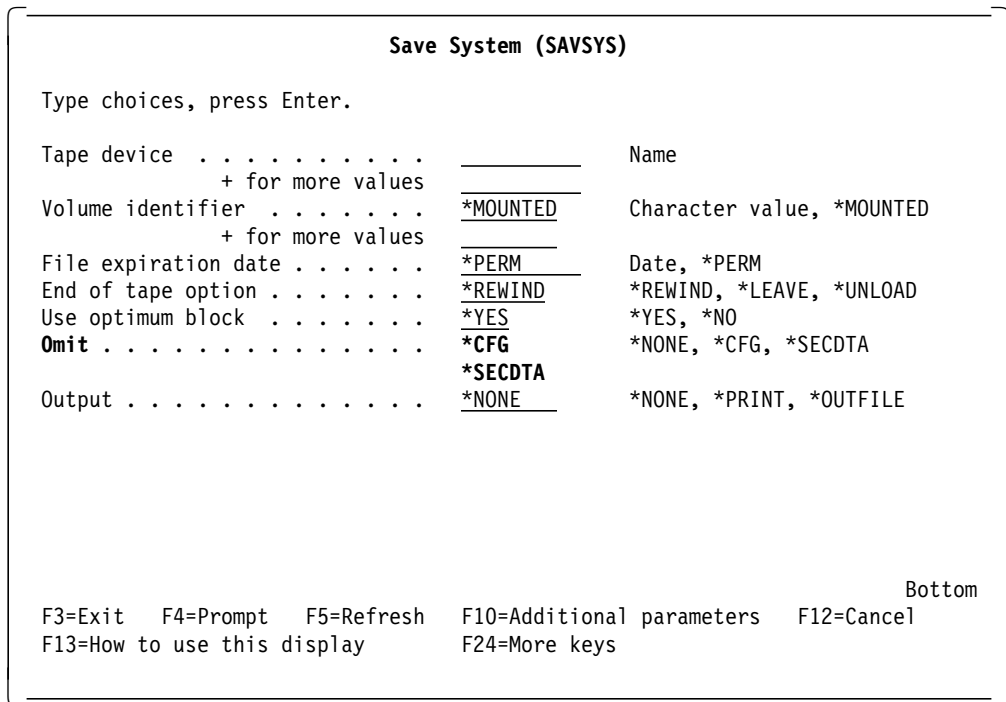


Figure 17. Save System Command

By omitting the configuration and security data from the SAVSYS operation, you reduce the amount of time that the system must be in a restricted state during the SAVSYS. In large environments, saving configuration objects and security information can take a long time.

Use the SAVCFG command instead of the SAVSYS command to save configuration data. Use the SAVSECDTA command instead of the SAVSYS command to save security data.

Note: The SAVCFG and SAVSECDTA commands do not require taking the system to a restricted state.

Recommendation

Perform both the SAVSECDTA and SAVCFG commands regularly to ensure that user profiles and configuration objects are saved.

Note: The OMIT parameter is supported on SAVSYS commands for OS/400 releases prior to V4R1 if the appropriate PTF is applied. The following PTFs or their supersedes contain this OMIT function for the SAVSYS command:

- V3R7—SF44526
- V3R6—SF38928
- V3R2—SF41420
- V3R1—SF44178

5.4 Concurrent Save Operations

Save operations on V4R2 systems are enhanced to shorten the time to save, allowing better flexibility while improving availability. Two changes are discussed in this section:

1. Concurrent saves for libraries
2. Concurrent saves for DLOs

5.4.1 Concurrent Saves on Libraries

On V4R1 and later, users can issue multiple SAVOBJ and SAVCHGOBJ commands, as well as the QSRSAVO API, against a single library at the same time. This allows users to issue multiple save operations and use multiple tape drives to save objects from a single large library. For example, you can save generic objects from a single large library to one tape drive and concurrently issue another SAVOBJ command against the same library to save a different set of generic objects on to another tape drive. This is due to the fact that the objects in the library are locked by the first save command.

The following example shows the concurrent SAVLIB request.

```
SAVOBJ OBJ(A* B* C* .... L*) LIB(MYLIB) DEV(TAP04)
SAVOBJ OBJ(M* N* O* .... Z*) LIB(MYLIB) DEV(TAP05)
```

In this example, objects starting with the letters A through L are saved on tape drive TAP04, while objects starting with the letters M through Z are saved on tape drive TAP05. Together the two SAVOBJ commands save all libraries beginning with an alphabetic character. (Remember to consider special characters such as #, *, or \$ when constructing your backup scheme).

Users can use all of the tape resources, and therefore, reduce the overall time to save. Installations with multiple tape drives and large libraries benefit most from this enhancement.

Note: If you are running concurrent saves of objects residing in the same library, tape drives do not need to be the same type. However, for ease of management and recovery, perform save operations to identical type tape drives.

You can also issue concurrent save commands against multiple libraries. When you run multiple save commands, the system processes the request in several stages that overlap, providing improved save performance. Figure 18 on page 56 shows how a save of multiple libraries occurs.

During the preprocessing phase for a library, the system creates a list of all the objects to be saved, called a save list, before the system starts actually copying data to the media. While data is copied to media during the postprocessing phase for that library, the system works on creating the save list for the next library or libraries to be processed.

SAVLIB LIB(LIBA LIBB LIBC LIBD)

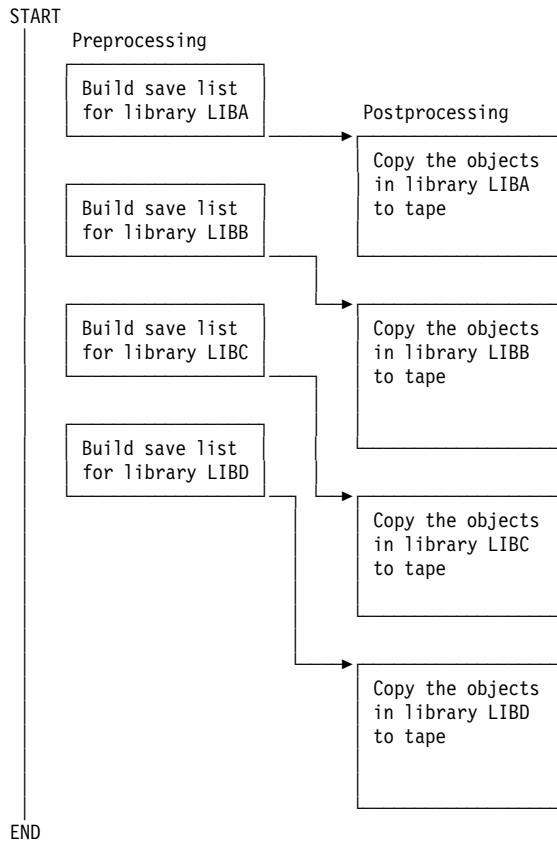


Figure 18. How the System Performs the Save

Refer to *Backup and Recovery*, SC41-5304, for more information.

5.4.2 Concurrent Saves for DLOs

On V4R1 and later, multiple SAVDLO operations can be performed concurrently for DLO objects within the same auxiliary storage pool (ASP). Issuing multiple SAVDLO commands enables the concurrent use of tape drives. To manage these multiple saves, the folder (FLR) parameter for the SAVDLO command supports generic folder names.

Note: If you run concurrent saves of DLOs, the tape drives do not need to be of the same type. However, for ease of management and recovery, consider performing save operations to identical type tape drives.

As with the SAVLIB and the SAVOBJ commands, you can specify up to 300 specific folder names or generic values for folders that are to be omitted from the SAVDLO operation. This enhancement allows users to omit folders that contain test data, object code that does not change frequently, or folders that are saved concurrently to a different tape drive. This is an advantage for installations that have archive folders or folders that do not change frequently.

With these changes, you can tailor your SAVDLO operation to take advantage of multiple tape drives and eliminate the need to save nonvolatile folder and document objects frequently. This reduces the time required for the SAVDLO operation and takes advantage of all of the tape drives that are installed on the system.

Here is an example of saving folders that are named generically:

```
SAVDLO DLO(*ALL)
      FLR(DEPT*)
      DEV(tape-device-name)
      OMITFLR(DEPT2* 'DEPT-A/WIN*')
```

In this example, all folders starting with “DEPT” are saved, with the exception of the folders that start with “DEPT2” and the sub-folders within folder DEPT-A that start with “WIN.”

Note: Omit folders are only allowed if DLO(*ALL) or DLO(*CHG) values are specified.

5.5 Concurrent Restore Operations

Concurrent restore operations have additional flexibility on V4R2 and later systems.

Users can also perform multiple concurrent restore operations across a single library, such as:

```
RSTOBJ OBJ(A* B* C* ...L*) SAVLIB(MYLIB) DEV(TAP04)
RSTOBJ OBJ(M* N* O* ...Z*) SAVLIB(MYLIB) DEV(TAP05)
```

In this example, objects starting with the letters A through L are restored from tape drive TAP04, while objects starting with the letters M through Z are restored from tape drive TAP05. This allows for a faster and more efficient recovery.

Similarly, users can perform multiple concurrent RSTDLO operations to a single ASP, using more than one tape device for faster recovery.

Note: If you implement an example based on the beginning letter of an object’s name, remember to consider objects that start with special characters, such as &, #, and so on.

5.6 Use Optimum Blocking for Save and Restore

Typically, tape drive performance relates to the speed of the device and to the block transfer size. Block transfer size is affected by the use optimum block (USEOPTBLK) parameter on the save commands.

Most of the throughput enhancements made in V3R1 (when the block size increased to 24KB) and V3R6 are realized with the 3590 and 3570 tape drive and the high-end processors. The time to transfer data to the drive is also improved significantly with the PowerPC AS processor hardware.

Beginning with V4R1, the USEOPTBLK parameter has been added to the SAVCFG, SAVSYS, SAVDLO, SAVSAVFDTA, and SAVSECDTA commands. This parameter was already enabled with V3R7 for the SAV, SAVOBJ, SAVCHGOBJ, and SAVLIB commands.

The default setting for USEOPTBLK was changed to *YES beginning with V4R2. With optimum block size set to *YES, blocking is enabled for the supported value (based on the device type) and is used on the save commands. If the block size that is used is larger than a block size that is supported by all device types, performance may improve, however:

- The tape file that is created is only compatible with a device that supports the block size used. Commands, such as Duplicate Tape (DUPTAP), do not duplicate files unless the files are duplicated to a device that supports the same block size that was used. You cannot duplicate a tape that is created using USEOPTBLK(*YES) to another tape drive that does not use USEOPTBLK(*YES).
- Data Compression (DTACPR) is not allowed if you use the optimum block size. If DTACPR(*DEV) is specified with USEOPTBLK(*YES) and the device supports an optimum block size, data compression is not performed.

The 3590 tape drive provides the best performance results when the USEOPTBLK parameter is set to *YES. The 3570 tape drive can also benefit from the new increased block size, but the performance gains are not as dramatic or as high as with the 3590.

For additional information on performance measurements, see *Backup and Recovery*, SC41-5304, and Appendix C, “Save and Restore Rates of IBM Tape Drives for Sample Workloads” on page 379.

Note: If the target release value specified in the save command is earlier than V3R7, the block size that is supported by all device types is used. That is, you cannot perform a save operation using an optimum block size to a release that does not support the USEOPTBLK parameter (releases prior to V3R7). For this reason, we recommend that you set the USEOPTBLK parameter to *NO if you plan to restore objects onto systems running releases prior to V3R7.

Starting with V4R2, the default setting of USEOPTBLK is *YES. This means that you can duplicate tapes from:

- 3590 to 3590
- 3570 to 3570
- 3590 to 3570
- 3570 to 3590

5.7 Save and Restore Scenario

We provide a save and restore scenario so you can better understand the impact of some of the options described in this chapter. The following example describes how elapsed save time is dramatically reduced by making use of various features of the save and restore functions, as well as with various hardware options. Our sample environment includes:

- **Hardware features**
 - 200G of DASD
 - Three 3590 tape drives
 - Each 3590 tape drive is connected to separate IOPs on separate buses
- **Software features**
 - OS/400 V4R2
 - System configuration with one ASP
 - 50 percent usage of the single ASP
- **Application and database features**
 - Two main production libraries

The following disk space report shows the disk usage and size of the main libraries on the system.

```

Disk Space Report
Page 4
5769SS1 V4R2M0 980228 SYSTEMXX 12/09/97 10:06:45

Library Information
Library Owner % of Size in Last Last Description
Disk 1000 bytes Change Use
PRODDATA QSYS 30.00 60000000.0 12/08/97 12/08/97 Production Lib 1
PRODDATA2 QSYS 10.00 20000000.0 12/08/97 12/08/97 Production Lib 2
QUSRSYS QSYS 3.25 6509872.6 12/08/97 12/08/97 System Library
QDOC QDOC 1.75 3507730.9 11/30/97 Document library
QFPINT QSYS .45 896288.5 10/27/97 12/07/97 Print file library
. . . . .
. . . . .
. . . . .
. . . . .
TOTAL 100000000.0

```

Figure 19. PRTDSKINF *LIB Listing

The disk space report (Figure 19) shows that the library, PRODDATA, represents 60 percent of the used storage space on SYSTEMXX. This is determined as follows:

$$\text{PRDDATA \% of Disk Utilization} / \text{Total DISK Utilization(Percentage)} = 30\% / 50\% = 60\%$$

Of the 60GB in the PRODDATA library, one half of the DASD usage is occupied by two files—CUSTMAST and SALESMAST.

```

Disk Space Report
Page 4
5769SS1 V4R2M0 980228 SYSTEMXX 12/10/97 12:15:48

Library and Objects Information
Library/ Object Type Owner % of Size in Last Last Description
Library 1000 bytes Change Use
PRODDATA *LIB QSYSOPR 60000000.0 12/08/97 11/05/97 PRODDATA production Library 1
CUSTMAST *FILE QSECOFR 33.34 20000000.0 12/08/97 11/05/97 Customer Master File
SALESMAST *FILE QSECOFR 16.67 10000000.0 12/08/97 11/03/97 Sales Master File
CSTMTR *FILE QSECOFR 10.00 6000000.0 10/24/97 11/03/97 B Customer File-has fields rearranged
HSTRY *FILE QSECOFR 8.00 4800000.0 10/24/97 B History File
ITEM *FILE QSECOFR 7.00 4200000.0 10/24/97 11/03/97 B Item File
ORDERS *FILE QSECOFR 7.00 4200000.0 10/24/97 11/03/97 B Orders
. . . . .
. . . . .
. . . . .
. . . . .
TOTAL 60000000.0
***** END OF LISTING *****

```

Figure 20. PRTDSKINF *LIB for Library PRODDATA

5.7.1 Save Strategy

In designing our save strategy, the first factor we look at relates to our 3590 hardware design. We have three 3590 tape drives connected to separate buses. Because the contention factor across these individual tape drives is removed by the fact they are on separate IOPs and separate expansion units, our first directive is to design a save strategy that uses all three tape drives concurrently.

To reduce the overall time to save required, we attempt to create our saves so that they begin and end together without complicating our strategy significantly.

We have 100GB that need to be saved. From the library information presented earlier, three saves are divided as follows:

- **Save 1:** SAVOBJ of the CUSTMAST and SALESMAST files in library PRODDATA

- **Save 2:** SAVLIB of library PRODDATA omitting the two objects CUSTMAST and SALESMAS saved separately in Save 1
- **Save 3:** An additional SAVLIB *ALLUSR omitting the library PRODDATA

By splitting the three saves, we save the following amount of space on each save:

- Save 1 = 30GB
- Save 2 = 30GB
- Save 3 = 40GB less the amount in SAVDLO and SAVSYS

The save command for the three saves in our design appear as shown here:

```
SAVOBJ OBJ(CUSTMAST SALESMAS) LIB(PRODDATA) DEV(3590TAP01)
```

```
SAVLIB LIB(PRODDATA) DEV(3590TAP02)
      OMITOBJ((PRODDATA/CUSTMAST) (PRODDATA/SALESMAS))
```

and

```
SAVLIB LIB(*ALLUSR) DEV(3590TAP03) OMITLIB(PRODDATA)
```

It is important to bring the system down as close to a restricted state as possible during this save for two reasons:

1. Other system resources used during the save process are minimized.
2. It ensures an accurate recovery of the system through a proper save of QUSRSYS.

If the save operation finds objects in use, message CPD3796 is issued on V4R2 and later to indicate:

```
. . Critical recovery data may not have been saved.
. . . Library QUSRSYS was not completely saved . . .
. . . . The most common reason why the objects were not saved
is because they were in use during the save operation.
These objects must be saved for a system recovery.
```

The user is referred to the job log for more information. By managing for this message, you can ensure that the QUSRSYS library is properly saved.

Recommendation

Monitor for the CPD3796 message and determine from the joblog what objects are not saved successfully. The named objects need to be saved from QUSRSYS along with the SAVSYS and SAVDLO performed when the system is in a restricted state.

To run the three saves in parallel when jobs are submitted concurrently, you can:

- Use the OS/400 job scheduler
- Set up control groups in BRMS
- Use CL programs

Or, use a combination of these options.

You must be in a restricted state to perform your SAVSYS. Also consider being in a restricted state to perform your SAVLIB LIB(*IBM), SAVDLO, and SAV of the integrated file system structures.

5.7.2 Additional Considerations

You may want to split up your incremental saves in the same fashion. The three incremental saves may appear similar to this:

- **Save 1:** SAVCHGOBJ OBJ(*ALL) LIB(*ALLUSR) DEV(3590TAP01)
OMITOBJ((PRODDATA/CUSTMAST) (PRODDATA/SALESMAST))
- **Save 2:** SAVOBJ OBJ(CUSTMAST) LIB(PRODDATA) DEV(3590TAP02)
- **Save 3:** SAVOBJ OBJ(SALESMAST) LIB(PRODDATA) DEV(3590TAP03)

Again, these saves can be performed through CL programs, BRMS, OS/400 job scheduler, or a combination of the three.

Note: When splitting up incremental saves, understand which objects in our application are being changed. The large files used in a full system backup to split the saves across the three tape drives may not necessarily be the files used in the split of your incremental saves. You need to track the changes to the large objects in your application libraries to properly divide your incremental saves efficiently.

Recommendation

The only way to ensure that you have a good backup strategy is to test your recovery. If you are not backing up all that you should, your test is done in a live recovery situation. This can produce undesirable results.

Contact your local service provider to find out the possibilities of testing your recovery procedures.

5.8 Unattended Saves Using the SAVE Menu

The SAVE menu has been enhanced over the last several releases to help the administrator manage backups. With the following changes, for example, fewer CL programs are required to manage the save operations. Changes include:

- SAVE menu additions for V4R2 systems:
 - The option to vary off the network servers (Integrated PC Server)
 - The option to unmount user-defined file systems (UDFS) before saving
 - The option to print system information along with the save (PRTSYSINF)

These additions are described in Section 5.8.2, “V4R2 Availability Options on the SAVE Menu” on page 63.

- SAVE menu additions for V3R1 and later:
 - Unattended saves using the Start time parameter

The following CL commands run when selecting options 21, 22, and 23 from the SAVE menu.

- Option 21 to Save Entire System runs these CL commands:

```
ENDSBS SBS(*ALL) OPTION(*IMMED)
SAVSYS
SAVLIB LIB(*NONSYS) ACCPTH(*YES)
SAVDLO DLO(*ALL) FLR(*ANY)
SAV OBJ(('/*) ('/QSYS.LIB' *OMIT)
(' /QDLS' *OMIT)) UPDHST(*YES)
STRSBS SBS(controlling-subsystem)
```

- Option 22 to Save System Data runs these commands:


```
ENDSBS SBS(*ALL) OPTION(*IMMED)
SAVSYS
SAVLIB LIB(*IBM) ACCPTH(*YES)
SAV OBJ('/QIBM/ProdData'
        ('/QOpenSys/QIBM/ProdData')) UPDHST(*YES)
STRSBS SBSD(controlling-subsystem)
```

- Option 23 to Save User Data runs these commands:


```
ENDSBS SBS(*ALL) OPTION(*IMMED)
SAVSECDTA
SAVCFG
SAVLIB LIB(*ALLUSR) ACCPTH(*YES)
SAVDLO DLO(*ALL) FLR(*ANY)
SAV OBJ('//*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)
        ('/QIBM/ProdData' *OMIT)
        ('/QOpenSys/QIBM/ProdData' *OMIT)) UPDHST(*YES)
STRSBS SBSD(controlling-subsystem)
```

Note that the ACCPTH parameter indicates *YES. This is the default when using option 21, 22, or 23 from the SAVE menu. ACCPTH(*YES) is *not* the default on save commands.

After selecting options 21, 22, or 23, a display appears that describes the function of the menu option that you selected. After reading the display and pressing Enter, the Specify Command Defaults display appears, as shown in Figure 21.

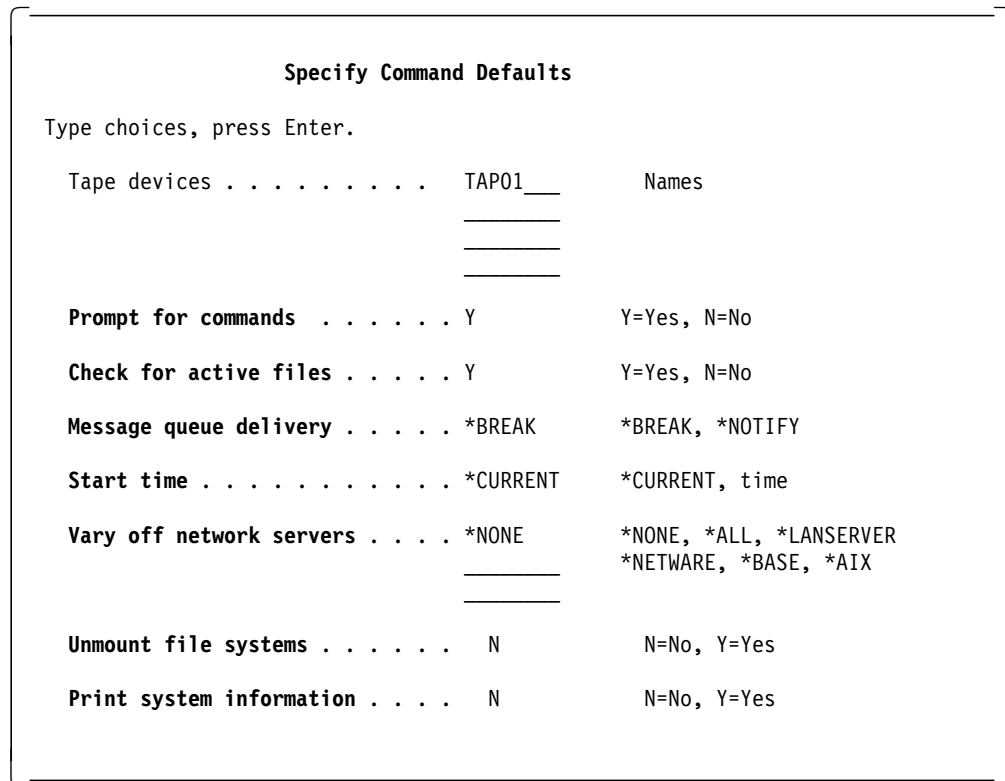


Figure 21. Specify Command Defaults

To set your saves for unattended mode with a delayed start, be sure to:

1. Change the prompt for commands to *N*.

2. Change the check for active files parameter to *N*.
3. Set up the system reply list.
 - Display the reply list sequence numbers to find what numbers are available for use. Use the command:
WRKRPYLE
 - If message CPA3708 is not represented in your reply list, add it. Use the command:
ADDRPYLE SEQNBR(nnnn) MSGID(CPA3708) RPY('G')

command. For nnnn, substitute an unused sequence number from 1 through 9999.

 - Change your job to use the reply list. Use the command
CHGJOB INQMSGRPY(*SYSRPLY)
4. Enter **NONE* for the QINACTITV system value.
5. Enter **NOTIFY* for message queue delivery.

If you enter **NOTIFY* for *Message queue delivery*, severity 99 messages not associated with the save operation are sent to the QSYSOPR message queue without interrupting the save process. This prevents communication messages to the operator from stopping the save operation.

5.8.1 Start Time

Specifying a start time in the HHMMSS or HH:MM:SS format, based on the 24-hour clock, allows you to schedule a save to start up to 24 hours after the command is submitted. However, there is a potential security risk of which you need to be aware. The workstation on which the user running the save using the SAVE menu remains signed on and waits for the jobs to start at the delayed time. If the system request function is used to cancel the job, the workstation returns to the SAVE menu. That is, it is not disconnected and the user remains signed on. This leaves the workstation signed on with whatever security authorization that user has, which may be **ALLOBJ* authority. To get around this, make sure that this workstation is in a secured environment.

5.8.2 V4R2 Availability Options on the SAVE Menu

On V4R2 and later, some SAVE menu parameters are added to allow you to:

- Indicate whether to vary off network servers before an option 21, 22, or 23 save is performed.
- Indicate whether to unmount user defined file systems before an option 21, 22, or 23 save is performed.
- Indicate whether to produce PRTSYSINF output.

5.8.2.1 Varying Off Network Servers

The vary off network servers parameter allows you to choose whether to vary off the network associated with the Integrated PC Server prior to the start of the save and vary on after completion of the save from the SAVE menu. If the network servers are varied off, the save contains a physical representation of the storage space associated with each logical PC drive, a disk image of the /QFPNWSSTG directories.

If the network servers are varied on, the save contains a logical representation of the data. For example, the /QLANSRV directories are associated with the Integrated PC Server.

Your choices include:

- *NONE

The network servers are not varied off. The save takes much longer because the data from all network servers are saved in a format that allows the restore of individual files and directories.

- *ALL

This option allows all network servers to be varied off prior to the start of the save. The save takes less time. Since the network servers storage spaces are saved with the network servers varied off, you cannot restore individual files or directories from this save.

- *LANSERVER

This option allows all network servers of type *LANSERVER to be varied off prior to the start of the save. The save takes less time, but you cannot restore individual files and directories.

- *NETWARE

This option allows all network servers of type *NETWARE to be varied off prior to the start of the save. The save takes less time, but you cannot restore individual files and directories.

- *BASE

This allows network servers of *BASE to be varied off prior to the start of the save. This option saves the network server storage spaces.

- *AIX

This option allows all network server of *AIX to be varied off prior to the save. It allows the save of the network server storage spaces.

On systems prior to V4R2, where the vary off network servers option is not available, vary off and on the network servers using the VRYCFG command.

5.8.2.2 Unmount File Systems

The unmount file system parameter allows you to indicate whether you want to have all dynamically mounted file systems unmounted prior to the start of the save.

If Y (for Yes) is selected, all dynamically mounted file systems are unmounted prior to the start of the save and all UDFSs and the objects contained in them are saved. The file systems are not remounted when the save completes. If N (for No) is selected, the save contains only path name information as if they are in the file system over the mounted file system. No information regarding the UDFS or ASPs containing the saved objects is saved. Objects contained in a directory, over which the UDFS is mounted, are not saved.

For example, directory /sxp1 contains objects and a UDFS is mounted over /sxp1. A save without unmounting the UDFS does not save objects in the /sxp1 directory.

Note: The User Defined File System (UDFS) must be unmounted to save the attributes associated with the file system. If the UDFS is mounted, the objects in the directory that are mounted over, are not saved. Only the objects in the UDFS are saved. The SAVE menu does not mount the file system back after the save.

Recommendation

For recovery purposes, unmount the UDFS prior to a save. Change the default from No to Yes.

5.8.2.3 Print System Information

The print system information parameter allows you to specify whether the system information is to be printed during the save. The system information report consists of lists containing how *your* system is configured. For example, it contains such information as:

- The libraries that are on the system and when they were last backed up
- How your system has implemented each system value
- A list of all system reply list entries
- A list of PTFs installed on the system
- Access path rebuild time estimates
- The power off and on schedule
- The output from DSPHDWRSC to document what hardware resources are installed with resource names identified
- A list of SNADS configuration objects
- A list of subsystem descriptions
- A DSPOBJD listing of system journal objects
- A DSPJRNA report
- A list of user profiles
- A list of job descriptions

5.9 ObjectConnect/400

ObjectConnect/400 is a set of six CL commands designed to move objects, libraries, or even integrated file system directories from one system to another without the need to implement a SNADS solution. Customers with more than one AS/400 system can use ObjectConnect/400 to:

- Create and maintain copies of critical objects, libraries, or integrated file system directories on other AS/400 systems for use during planned outages in addition to disaster recovery. ObjectConnect/400 also allows the copying of the objects back to the original AS/400 system after an outage.
- Migrate objects including document library objects (DLOs), folders, libraries, or integrated file system directories from one AS/400 system to another during system upgrades without disruption.
- Distribute objects, libraries, or integrated file system directories to other AS/400 systems in a network, allowing other systems to efficiently refer to local copies of the information.

Beginning with V3R6 RISC and V3R2 CISC systems, ObjectConnect/400 is included in OS/400 as a no-charge feature. For V3R1 systems, the support is available with a PRPQ (number P84244).

ObjectConnect/400 is installed as option 22 of the base operating system for V3R2 and V3R6. You can use the Display Software Resource (DSPSFWRSC) command or option 10 on the LICPGM menu to view licensed programs and confirm ObjectConnect/400 installation. ObjectConnect/400 objects are stored in the library QSR. The CL commands are in library QSYS.

5.9.1 ObjectConnect/400 Command Sets

Table 2 shows the six ObjectConnect commands.

<i>Table 2. ObjectConnect/400 and Equivalent AS/400 Commands</i>	
ObjectConnect/400 Save and Restore Commands	Equivalent AS/400 Save and Restore Commands
SAVRST Save/Restore Integrated File System	Save (SAV), Restore (RST)
SAVRSTOBJ Save/Restore Object	Save Object (SAVOBJ), Restore Object (RSTOBJ)
SAVRSTCHG Save/Restore Changed Objects	Save Changed Objects (SAVCHGOBJ), Restore Object (RSTOBJ)
SAVRSTLIB Save/Restore Library	Save Library (SAVLIB), Restore Library (RSTLIB)
SAVRSTDLO Save/Restore Document Library Objects	Save Document Library Objects (SAVDLO), Restore Document Library Objects (RSTDLO)
SAVRSTCFG Save/Restore Configuration	Save Configuration (SAVCFG), Restore Configuration (RSTCFG)

ObjectConnect/400 CL commands provide the same function as standard save and restore commands. These functions include: save-while-active support, CISC-to-RISC migrations (forced object conversions), and previous release support. ObjectConnect/400 save and restore commands support save and restore functions the same as on standard AS/400 Save and Restore commands. These functions include:

- Authority and security checking
- File and member control
- Release-to-release compatibility
- Object locks considerations
- Allow object differences
- Storage freed
- ASP number

The exceptions to this are SAVRSTCFG, which does not support as many object types as the SAVCFG (save configuration) and RSTCFG (restore configuration) commands. The list of configuration objects supported by SAVRSTCFG includes:

*CFGL	*DEVD	*NTBD
*CNL	*IPXD	*NWID
*COSD	*LIND	*NWSD
*CTLD	*MODD	*SRM

Note: If you use OPTION(*NEW) for ObjectConnect/400 command parameters on a large library or object, ObjectConnect/400 saves all the objects, sends them over to the second system, but only restores what is required. So if the data is large, it is possible that a lot of unnecessary processing is performed.

You can access the ObjectConnect/400 commands from the CMDSAVRST menu or enter them directly on the command line.

If you want to save and restore object X from System A to System B, ObjectConnect/400 performs this operation using a communications link between the two systems. If you want to save and restore object Y from System A to System C and you have Opticonnect/400 operating between these systems, Object Y is restored to System C using the fiber bus. If you have both Opticonnect/400 and a communications link between the two systems, ObjectConnect/400 tries to use the faster fiber optic bus connection first. If that connection cannot be made, the operation automatically uses an APPC controller definition instead. Note that if an OptiConnect/400 fiber bus is available, it is not necessary to have a communications link for ObjectConnect/400.

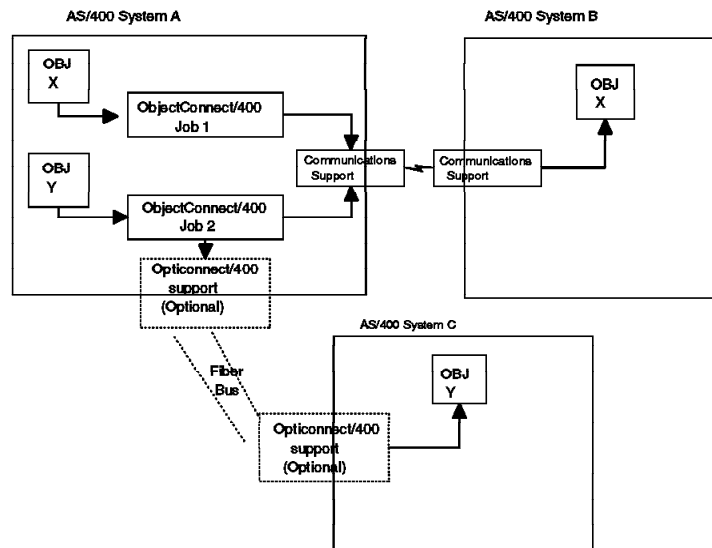


Figure 22. ObjectConnect/400 using Opticonnect/400 or Communications Link

ObjectConnect/400 complements existing distribution techniques, such as SNADS and tape and in some cases may be a better alternative. Because ObjectConnect/400 does not use internal distribution queues or save files, it can be more efficient and easier to manage than SNADS. And because it is a single AS/400 CL command, such as SAVRST, it lends itself to an easier automation path.

5.9.2 ObjectConnect/400 Offers Simplicity

One ObjectConnect/400 CL command provides both saves and restores, for example:

```
SAVRSTLIB LIB(SXP) SYSTEM(SYSTEMXX)
```

Figure 23 on page 68 shows how using ObjectConnect/400 results in a reduction in I/O and system resources when compared to SNADS.

SNADS distribution of a library or object requires the use of a save file, and in turn, is copied to a distribution queue before it is sent. At this stage, we have two additional copies of the object and library—one copy using the CRTSAVF command and the other using the SNDNETF command. This results in additional disk space and system and I/O resources to manage this activity.

These same issues are replicated on the target AS/400 system. A temporary save file is created to receive the save file from the distribution queue.

ObjectConnect/400 uses up to two-thirds less DASD resources than SNADS for this scenario. As soon as the ObjectConnect/400 command begins saving on the source AS/400 system, data is transmitted and the restore process starts on the target AS/400 system.

In comparing ObjectConnect/400 with other methods of moving objects, such as tape units, save files, and SNADS (as just described in this section), there are some obvious advantages to ObjectConnect/400.

1. ObjectConnect/400 uses a CL command interface designed for those familiar with save and restore commands on the AS/400 system. To move an object from one system to another, you enter one CL command:

```
SAVRSTLIB LIB(DMT) SYSTEM(SYSTEMXX)
```

2. ObjectConnect/400 offers efficiency advantages in both processor usage and the amount of disk storage used. ObjectConnect/400 operates using interfaces that reduce required system resources on both the source and target systems.
3. ObjectConnect/400 does not use intermediate files, such as the use of save files in a SNADS implementation. This reduces processor and disk usage.

Figure 23 shows how the ObjectConnect/400 and SNADS solutions duplicate a library and send it across to another system.

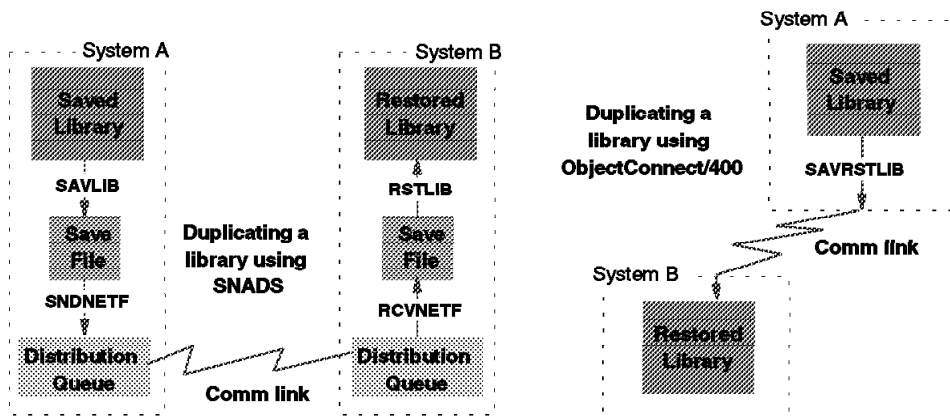


Figure 23. Duplicating a Library Using SNADS versus ObjectConnect/400

To summarize, ObjectConnect/400 advantages include:

- ObjectConnect/400 is ideal for repetitive tasks that require automation.
- Since ObjectConnect/400 is not limited to a single library during its restore process, it deals with multi-library distribution easier than tape, save files using tape, or save files through SNADS.
- If an AS/400 site has limited tape handling or operator resources, using ObjectConnect/400 can increase productivity by using ObjectConnect/400 for library and object distribution.

- Moving applications from one AS/400 system to another is easier without the need to use tape media that reduces the reliance on an availability of tape devices.
- ObjectConnect/400 helps eliminate an incompatibility in tape devices between AS/400 systems.

5.9.3 ObjectConnect/400 Implementation Considerations

Here is a list of considerations when implementing an ObjectConnect/400 solution for your system:

- ObjectConnect/400 is not a replacement for high availability solutions, such as the applications developed by high availability vendors. High availability solutions include dual systems to provide hot site backup capabilities, recovery automation, and data synchronization functions. ObjectConnect/400 only provides periodic data duplication. However, ObjectConnect/400 is useful for implementing and maintaining a dual-system solution, providing initial copying of data, and performing periodic refreshes if needed.
- ObjectConnect/400 does not replace a multi-system database function, such as DataPropagator.
- ObjectConnect/400 does not replace a requirement for a valid backup and recovery plan, which includes a primary backup on tape.
- ObjectConnect/400 does not replace the need for tape backups to be stored offsite.

Note: ObjectConnect/400 does not issue status messages as the equivalent save and restore commands do. Status messages are useful to watch the progress of the restore, as objects are not named during the ObjectConnect/400 restore operation.

Refer to Appendix E, “High Availability Solutions” on page 391 for high availability solutions and Chapter 18, “OptiConnect for OS/400” on page 351.

5.10 Save and Restore Spooled Files

You can save and restore spooled files using the following methods (apart from the BRMS product):

1. Use the save and restore spooled file (SAVRSTSPLF) tool found in the QUSRTOOL library
2. Create your own commands to save and restore spooled files.

These two methods are described in the following sections.

5.10.1 QUSRTOOL to Save Spooled Files

The SAVRSTSPLF tool consists of the ZSAVSPLF and ZRSTSPLF programs. This tool provides a method for saving spooled files into a designated library to restore and print later. All information necessary to use this tool is documented in the TSRINFO member of the QATTINFO file in the QUSRTOOL library. QUSRTOOL is optionally installed as the example tools library option 7 of OS/400.

The ZSAVSPLF command stores the files specified by the user to a designated library or device. On the ZSAVSPLF command, the spooled files to be saved can be selected by user, output queue, form type, or user data.

The ZRSTSPFL command restores those files from the library or device to the output queue in which they were originally spooled.

Note the following considerations of the SAVRSTSPFL tool:

- Only spooled files intended for printing can be saved. Diskette spooled files are skipped.
- User space name generation for each spooled file uses the first six characters of the job name and a four digit number. This limits the number of spooled files saved to 9 999.
- Any error condition reported to the save and restore program terminates the operation.
- The largest spooled file cannot exceed 16MB.
- The original owners user profile must be on the system to which the spooled file is being restored.

5.10.2 Creating CL Commands to Save Spooled Files

You can also create commands to save and restore spooled files. The following is a list of instructions for creating and using commands that we name GETSPLF and PUTSPLF. If you create the programs on your own system, use names to fit your naming convention.

5.10.2.1 Creating a Get Spooled File (GETSPLF) Command

1. To create the GETSPLF CL command on your system, type:

```
CRTSRCPF FILE(QUSRSYS/QCMDSRC) TEXT('Command Source File')
```

2. Type the following into the source file, QUSRSYS/QCMDSRC with member name GETSPLF:

```
GETSPLF:CMD  PROMPT('Get Spooled File')
             PARM      KWD(FILE)  +
                   TYPE(*NAME) LEN(10) RTNVAL(*NO)  +
                   RSTD(*NO) MIN(1) MAX(1)  +
                   FILE(*IN) FULL(*NO) EXPR(*YES) VARY(*NO)  +
                   PASSATR(*NO)  +
                   PROMPT('Spooled file')
             PARM      KWD(TOFILE)  +
                   TYPE(Q1) RTNVAL(*NO) MIN(1) MAX(1)  +
                   FILE(*OUT) PROMPT('To data base file')
             PARM      KWD(JOB)  +
                   TYPE(Q2) RTNVAL(*NO) DFT(*)  +
                   SNGVAL(*) MIN(0) MAX(1)  +
                   FILE(*NO) PROMPT('Job name')
             PARM      KWD(SPLNBR)  +
                   TYPE(*INT2) RTNVAL(*NO)  +
                   RSTD(*NO) DFT(*ONLY) RANGE(1 9999)  +
                   SPCVAL((*ONLY 0) (*LAST -1))  +
                   MIN(0) MAX(1)  +
                   EXPR(*YES) VARY(*NO) PASSATR(*NO)  +
                   PROMPT('Spooled file number')
             PARM      KWD(TOMBR)  +
                   TYPE(*NAME) LEN(10) RTNVAL(*NO)  +
                   RSTD(*NO) DFT(*FIRST) SPCVAL(*FIRST)  +
                   MIN(0) MAX(1) FILE(*NO)  +
                   FULL(*NO) EXPR(*YES) VARY(*NO)  +
                   PASSATR(*NO) PROMPT('To member')
Q1:          QUAL  TYPE(*NAME)  +
             LEN(10)  +
```



```

RSTD(*NO) +
MIN(1) +
FULL(*NO) +
VARY(*NO) +
EXPR(*YES) +
PASSATR(*NO)
QUAL TYPE(*NAME) +
LEN(10) +
RSTD(*NO) +
DFT(*LIBL) +
SPCVAL((*LIBL (*CURLIB *CURLIB)) /* */ +
MIN(0) FULL(*NO) +
VARY(*NO) +
EXPR(*YES) +
PASSATR(*NO) +
PROMPT(' Library')
1Q2: QUAL TYPE(*NAME) +
LEN(10) +
RSTD(*NO) +
MIN(1) +
FULL(*NO) +
VARY(*NO) +
EXPR(*YES) +
PASSATR(*NO)
QUAL TYPE(*NAME) +
LEN(10) +
RSTD(*NO) +
MIN(0) FULL(*NO) +
VARY(*NO) +
EXPR(*YES) +
PASSATR(*NO) +
PROMPT(' User')
QUAL TYPE(*CHAR) +
LEN(6) +
RSTD(*NO) +
RANGE(000000 999999) +
MIN(0) FULL(*YES) +
EXPR(*YES) +
PASSATR(*NO) +
PROMPT(' Number')

DLTCMD CMD(QUSRSYS/GETSPLF)
CRTCMD CMD(QUSRSYS/GETSPLF) PGM(QUSRSYS/QSPGETF) +
SRCFILE(QUSRSYS/QCMDSRC) TEXT(' Command Source File')

```

3. Compile the source and correct any compile errors.

5.10.2.2 Create a Put Spooled File (PUTSPLF) Command

1. To create the PUTSPLF CL command on your system type:

```
CRTSRCPF FILE(QUSRSYS/QCMDSRC) TEXT(' Command Source File')
```

- Type the following into member PUTSPLF in the QUSRSYS/QCMDSRC source file:

```

PUTSPLF: CMD PROMPT(' Put Spooled File')
PARM KWD(FROMFILE) +
TYPE(Q1) MIN(1) MAX(1) +
FILE(*IN) PROMPT(' From file')
PARM KWD(OUTQ) +
TYPE(Q1) MIN(1) MAX(1) +

```

```

        FILE(*NO) PROMPT(' Output queue')
    PARM  KWD(FROMMBR)  +
        TYPE(*NAME) LEN(10)  +
        DFT(*FIRST) SPCVAL(*FIRST)  +
        MIN(0) MAX(1) FILE(*NO)  +
        FULL(*NO) EXPR(*YES) VARY(*NO)  +
        PASSATR(*NO) PROMPT(' From member')
Q1:     QUAL TYPE(*NAME)  +
        LEN(10)  +
        RSTD(*NO)  +
        MIN(1)  +
        FULL(*NO)  +
        VARY(*NO)  +
        EXPR(*YES)  +
        PASSATR(*NO)
        QUAL TYPE(*NAME)  +
        LEN(10)  +
        RSTD(*NO)  +
        DFT(*LIBL)  +
        SPCVAL((*LIBL) (*CURLIB *CURLIB)) /* */ +
        MIN(0) FULL(*NO)  +
        VARY(*NO)  +
        EXPR(*YES)  +
        PASSATR(*NO)  +
        PROMPT(' Library')

```

```

DLTCMD CMD(QUSRSYS/PUTSPLF)
CRTCMD CMD(QUSRSYS/PUTSPLF) PGM(QUSRSYS/QSPPUTF)
SRCFILE(QUSRSYS/QCMDSRC) TEXT(' Put Spooled File')

```

2. Compile the source and correct any resulting errors.

When using these CL commands, consider these points:

- They copy only one spooled file at a time.
- They copy spooled files to a database file member only.
- When the spooled file is restored, the person who runs the PUTSPLF command becomes the owner of that spooled file.

5.11 Multinational Environments and Object Names

For those of you who operate in a multilingual system environment or use multilingual support, there are additional considerations to ensure objects are usable on all systems in the business network. If you operate in mixed language environments, you need to understand the implications of using “special” characters for a naming convention.

Even if the names work perfectly on your local system and you can save and restore without problems, there can be problems when restoring on a system that uses a different character set ID (CCSID). Problems can arise when a special character is restored on an American system, such as when a Danish AA (which is actually printed as an A with a circle on top) is converted into a \$ sign on an American system.

Recommendation

Do not use special characters in:

- Library names
- File names
- File member names

See the redbook *Speak the Right Language with Your AS/400 System*, SG24-2154, to learn more about language sensitive settings.

5.12 Product Preview for Save Restore Enhancements

Hardware disk compression using the LZ1 compression algorithm is present in disk IOPs available on V4R2 systems to manage the movement and placement of save and restore data. IBM intends to provide an update of OS/400 that will deliver Hierarchical Storage Management (HSM) and disk compression support for the following disk controllers:

- Feature code 2741 on Models 620 and S20
- Feature code 6533 on Models 620, 640, 650, S20, S30, S40, SB1, 500, 510, 530, 50S, and 53S
- Feature code 9754 on Models 640, 650, S30, S40, and SB1

This combination of compressed DASD with supporting software will offer customers save and restore efficiencies to more efficiently handle large volumes of data.

You will find additional information on HSM and disk compression at product announcement time, or contact your local IBM representative in the interim.

Chapter 6. Save and Restore Considerations for Mixed Release Environments

Many customers who maintain multiple AS/400 systems routinely save objects from one system and restore them on another. Release-to-release support on the AS/400 system allows you to copy objects from a current release system to a previous release system. This support also allows you to copy objects from a previous release system to a current release system.

Customers with more than one AS/400 system in their environment often operate (temporarily or permanently) different releases of OS/400 on different AS/400 systems in their environment. In many cases, customers maintain a previous release system as part of a high availability implementation. In this case, running different releases is part of the availability plan. Some customers maintain mixed releases when staging an upgrade in a 24 x 7 clustered system environment since not all systems can be unavailable at the same time.

The mix of different operating system levels requires the system administrator to make extra considerations in managing availability and recovery. For example, in some cases, the administrator must take production databases offline for several hours to perform a save on the current system, and subsequently restore to a previous release system. The administrator must know whether the restored objects are compatible with the target system.

Performance of the save and restore operations, disk use, or the capacity to store objects are also considerations.

This chapter discusses save and restore considerations for managing AS/400 systems at mixed release levels to increase reliability and flexibility. It also addresses:

- Considerations for using the USEOPTBLK parameter when the target system is at a previous release
- Observability considerations for previous release target systems

6.1 Target Release and Save While Active

On systems prior to V4R2, an object cannot be in use when saving it for use on a previous release system. That means, the save while active (SAVACT) parameter can only specify a *NO value when the target release (TGTRLS) parameter indicates a release prior to that of the system where the object is saved. If this combination is attempted, the save command fails with a CPD3728 diagnostic message indicating:

Specified TGTRLS value not allowed with SAVACT value...

In other words, TGTRLS has to indicate *NO. As the allowed window of time in which the operation has to perform saves decreases, insurance of a successful save becomes more difficult to administer when the object is targeted for another release level system.

On V4R2 systems, the object is saved for use on a previous release system *even if it is in use* during the save operation. Any valid target release (TGTRLS) value

can be specified on the save commands that allow the Save While Active (SAVACT) function.

Refer to Table 3 for a list of the supported target releases.

Current-release-to-previous-release support enables objects that are created and saved on the current release (V4R2) to be restored and used on any supported previous release. Most object types are supported on both release levels as long as they use functions from a previous release only.

Note: CRTxxxPGM commands are also affected by TGTRLS considerations. Be aware that specifying only TGTRLS on save commands does not ensure object compatibility. You must specify the TGTRLS parameter when the object is created. If not, the object is not saved, and a message is logged in the job log alerting the operator to that condition.

Recommendation

Use the TGTRLS parameter on save *and* create commands if the restore is to be done on a previous release system. This ensures that the object can be restored in a format that is usable on the intended target system.

Table 3 illustrates which TGTRLS values can be used on the create and save commands to prepare the object for restoring on a previous (or current) release system. The column *Other Valid *PRV Values* shows to which release level systems an object can be restored.

Source OS/400 Release	*CURRENT	*PRV	Other Valid *PRV Values
V4R2M0	V4R2M0	V4R1M0	V3R2M0 V3R7M0 V4R1M0 V4R2M0
V4R1M0	V4R1M0	V3R7M0	V3R1M0 V3R2M0 V3R6M0 V3R7M0 V4R1M0
V3R7M0	V3R7M0	V3R6M0	V3R0M5 V3R1M0 V3R2M0 V3R6M0 V3R7M0
V3R6M0	V3R6M0	V3R1M0	V2R3M0 V3R0M5 V3R1M0 V3R6M0
V3R2M0	V3R2M0	V3R1M0	V2R3M0 V3R0M5 V3R1M0 V3R2M0
V3R1M0	V3R1M0	V2R3M0	V2R3M0 V3R0M5 V3R1M0

On V4R2 systems, the target release (TGTRLS) parameter is allowed on these five OS/400 commands:

- Save Changed Objects (SAVCHGOBJ)
- Save Library (SAVLIB)
- Save Object (SAVOBJ)
- Save Object (SAV) — intended for use to save integrated file system objects
- Save Document Library Objects (SAVDLO)

The TGTRLS parameter is also indicated on the QSRSAVO API, and on these ObjectConnect/400 commands:

- Save Restore Library (SAVRSTLIB)
- Save Restore Object (SAVRSTOBJ)
- Save Restore Object (SAVRSTOBJ)
- Save Restore Changed Objects (SAVRSTCHG)
- Save Restore Doc/Lib Objects (SAVRSTDLO)
- Save Restore (SAVRST)

Note: TGTRLS support that uses Save While Active for save commands on V4R1, V3R7, V3R6, and earlier is restricted to database objects only. Support for the additional target releases is added by applying PTFs containing the updated function. Refer to Table 4 for the applicable PTF for your system.

OS/400 Release	PTF Number	APAR Number
V4R1	SF41474	SA63980
V3R7	SF38191	SA59311
V3R6	SF38229	SA59311

Contact the IBM Support Line or refer to the applicable APAR mentioned in Table 4 for more information. To order PTFs or APARs, use SNDPTFORD. See Chapter 11, "Availability and the PTF Process" on page 167 or the manual *Basic System Operation, Administration, and Problem Handling*, SC41-5206, for information on ordering PTFs.

For more information about release-to-release support, refer to *Backup and Recovery*, SC41-5304.

6.2 Observability Considerations when Restoring Objects

Observability is an important consideration when restoring a program from a previous release onto a newer OS/400 release system. To migrate programs to newer releases, the programs must have observability.

An observable object is one that has a program template associated with it. The program template contains information external to the execution of the program, such as a list of statement addresses and addresses used by program variables. The restore function uses the template to re-encapsulate the object into the new release. If the object is observable, the AS/400 system resolves the object during the restore function to re-encapsulate it into the new release. If programs have observability, they are migrated without recompiling. If programs have no observability (no template), they must be recompiled on the AS/400 system from a source member.

The program template allows the user to trace, set breakpoints, and observe variables when running the program in debug mode. These functions are not available if the program template is not present when the restore function re-encapsulates the object into the new release.

Note: To regain the program template (that is, to regain observability), the program must be recompiled from the associated source member.

To see if a program is observable, issue one of these commands:

```
DSPPGM
```

or

```
CHGPGM FRCCRT(*YES)
```

Note: CHGPGM forces a re-creation of the program. Use it with caution since it can be a long running process.

It is typical for observability to be retained at the central site AS/400 system and removed for copies sent to other systems in the network. Although the programs execute without observability, program debug capabilities are not available. Since the object is smaller in size, less resource is required to transport the object to the previous release system. Removing observability is also useful when the amount of disk space is limited.

To remove observability from OPM and ILE program objects, enter:

```
CHGPGM PGM(program-name) RMVOBS(*ALL) for OPM programs, or
```

```
CHGMOD MODULE(module-name) RMVOBS(*ALL) for ILE modules.
```

Note: In general, do not remove observability from programs that you intend to save on a RISC-based system (running V3R6 or later) and restore to a CISC-based system (V3R2 or earlier).

For more information on removing observability and the CHGPGM command, see the *CL Reference*, SC41-5722.

6.3 USEOPTBLK for Save Performance

The Use Optimum Block (USEOPTBLK) parameter allows you to indicate that you want data written to tape using the optimum block size supported by that particular tape drive. By specifying a *YES value on the USEOPTBLK parameter, performance improvements up to 50% are realized for the tape drives that support optimum blocking (such as the 3590, 3570 and 7208 Model 342).

With USEOPTBLK set to *YES, you make optimal use of any tape drive. With *YES specified for a tape device that does not support optimum blocking (for example the 9348 and 6385), the parameter is essentially ignored. There is no corresponding performance gain nor loss.

On V3R7 systems (and not on prior systems), the USEOPTBLK parameter is available on these save commands:

- SAV (Save)
- SAVLIB (Save Library)
- SAVCHGOBG (Save Changed Objects)
- SAVOBJ (Save Object)

In V4R1 and later, the USEOPTBLK support is available on these additional save commands:

- SAVSYS (Save System)
- SAVCFG (Save Configuration)
- SAVSECDTA (Save Security Data)
- SAVDLO (Save Document Library Object)
- SAVSAVFDTA (Save Save File Data)

It is also available on the V3R7 QRSAAVO API.

6.3.1 USEOPTBLK and TGTRLS Compatibility

You can combine the TGTRLS and USEOPTBLK parameters to maintain a previous release system, but only if the target release system tape device supports optimum blocking.

```

                                Save Library (SAVLIB)
Type choices, press Enter.
Library . . . . . > EDALIB      Name, generic*, *NONSYS...
      + for more values
Device . . . . . > TAP01       Name, *SAVF
      + for more values
      + for more values
Sequence number . . . . . *END   1-16777215, *END
Label . . . . . *LIB
File expiration date . . . . . *PERM   Date, *PERM
End of tape option . . . . . *REWIND  *REWIND, *LEAVE, *UNLOAD
Use optimum block . . . . . > *YES   *NO, *YES
Target release. . . . . . > V3R7MO *CURRENT, *PRV, V3R2MO
Update history . . . . . *YES      *YES, *NO
  
```

Figure 24. USEOPTBLK and TGTRLS Compatibility

If the TGTRLS value indicates a release that does not support optimal block size, the save operation completes, but the optimal block size is not used (the parameter is ignored). On systems prior to V4R2, this same situation results in a CPD378A message indicating:

Parameters not valid with USEOPTBLK value. When the use optimum block (USEOPTBLK) parameter value is specified as *YES, the target release (TGTRLS) parameter value cannot be specified as a release earlier than V3R7. Also, the target release of the save file cannot be a release earlier than V3R7.

If the TGTRLS value specified is earlier than V3R7, a block size supported by all device types is used.

The earliest target release that supports USEOPTBLK is depicted in the following table.

Table 5. Earliest Target Release	
Command	Earliest Target Release
SAVLIB	V3R7
SAVOBJ	V3R7
SAVCHGOBJ	V3R7
SAV	V3R7
SAVDLO	V4R1
SAVSYS	*CURRENT

Note: The SAVSYS, SAVSECDTA and SAVCFG commands are not used for restoring objects to a previous release system, and therefore, do not support the TGTRLS parameter.

Note: The SAV, SAVOBJ, SAVCHGOBJ, and SAVLIB commands have supported optimum blocking since V3R7. Optimum blocking support was added for the

SAVCFG, SAVSYS, SAVDLO, SAVSAVFDTA, and SAVSECDTA commands in V4R1.

6.3.2 USEOPTBLK and DTACPR Compatibility

You can combine the data compression (DTACPR) and USEOPTBLK parameters. If DTACPR(*YES) is specified and the optimal block size is 32K or larger, data compression is not performed and optimal block is used. If the optimal block size is 32K or smaller, data compression is performed. On systems prior to V4R2, this same situation results in message CPD378A indicating:

Parameters not valid with USEOPTBLK value...

When the use optimum block (USEOPTBLK) parameter value is specified as *YES, the data compression (DTACPR) parameter value cannot be specified as *YES, . . .

If DTACPR(*YES) is specified with a TGTRLS value that does not support optimal block size, the save completes, and data compression is performed. However, optimal block size is not used. On systems prior to V4R2, this same situation results in message CPD378A indicating:

Parameters not valid with USEOPTBLK value...

When the use optimum block (USEOPTBLK) parameter value is specified as *YES, the data compression (DTACPR) parameter value cannot be specified as *YES, the target release (TGTRLS) parameter value cannot be specified as a release earlier than V3R7M0, or the target release of the save file cannot be a release earlier than V3R7M0.

6.3.3 USEOPTBLK and Other Considerations

When using USEOPTBLK, consider the following points:

- When the QlpHandleCdState API is used to put the system into a CD-ROM state before performing a save, the save completes using a block size that is compatible with CD-ROM mastering. On systems prior to V4R2, the same situation results in a CPF384E message indicating:
USEOPTBLK(*YES) not valid for CD-ROM premastering...
- Tapes produced are only duplicated to a tape drive that supports the same block size using the Duplicate Tape (DUPTAP) command.

Recommendation

If a mix of release levels and tape device types are used for data exchange, specify USEOPTBLK(*NO) to ensure compatibility for the Duplicate Tape (DUPTAP) and save and restore functions.

- If USEOPTBLK(*YES) is specified when DTACPR(*YES) is also specified, the tape is written without data compression, and information message CPI3818 is logged in the job log indicating that data compression is not performed when optimum block size is specified. On systems prior to V4R2, users receive message CPD378A indicating:

Parameters not valid with USEOPTBLK value. When the use optimum block (USEOPTBLK) parameter value is specified as *YES, the target release (TGTRLS) parameter value cannot be specified as a release earlier than V3R7, or the target release of the save file cannot be a release earlier than V3R7.

That is, the DTACPR is ignored.

- If USEOPTBLK(*YES) is specified on the SAVSAVFDTA command and the save file is saved for a target release that does not support optimal block sizes, the tape is written with a block size that is compatible with the specified release. On systems prior to V4R2, the user receives the message CPD378A.
- When the user saves to a save file, no error message appears. On systems prior to V4R2, the user receives the diagnostic message CPD3754 indicating:
Save file and device parameters cannot be used together.
- When the user is using an optical device, no error message appears. On systems prior to V4R2, the user receives a CPD376E diagnostic message indicating:
Parameters not valid with optical device . . .

6.3.4 Correcting a Back Level QUSRSYS and QGPL Library on a CISC to RISC Migration

A mismatch can occur between the QGPL and QUSRSYS libraries as compared to other IBM-supplied libraries on the system when a conversion is done from a CISC release to a RISC system using the RSTLIB command. The RSTLIB command overlays the target system (RISC release) with objects from the source system (CISC release).

This presents a problem since the IBM-supplied objects in the QGPL and QUSRSYS libraries are not compatible with the system objects in other libraries. Although QGPL and QUSRSYS can contain user objects, these libraries cannot be restored without considering how to prevent user objects from being overlaid.

In this situation, it is possible to reconcile the QGPL and QUSRSYS library release levels without a complete system restore. The process assumes that you are familiar with central site distribution of tapes and DSLO media, as described in *Central Site Distribution*, SC41-5308.

The following steps outline what you can do to correct a mismatch of the QGPL and QUSRSYS libraries on a CISC to RISC migration.

1. Make sure you have a save of QGPL and QUSRSYS from the CISC system.
2. Create DSLO tapes from current RISC system. Use the *Central Site Distribution*, SC41-5308. to create the DSLO tapes.
(Hint: Use option 40 from the Work with Licensed Programs menu.)
The DSLO process asks if you want to create an installation profile. Use command key F3 to bypass this prompt.
3. When you are prompted to enter the SAVSYS command, press ENTER.
4. At the SAVLIB LIB(QGPL) prompt, press ENTER.
5. When the SAVLIB LIB(QUSRSYS) prompt appears, press ENTER.
6. The same screen appears again as an option 13 GO LICPGM screen. Type a "1" beside *all* the licensed program products in the list. Press ENTER.
7. On the RISC system, enter:

```

SAVLIB LIB(*ALLUSR) DEV(tape-device-name)
OMITLIB(QGPL QUSRSYS) ACCPTH(*YES)
SAVDLO DLO(*ALL) DEV(tape-device-name)
SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ(('/*')+
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))

```

8. On the RISC system, save the QGPL and QUSRSYS libraries onto a separate set of tapes. This tape set is used to restore user created objects.

9. Scratch install using the tape created in step 2.

Note: Refer to the appropriate section of *Backup and Recovery*, SC41-5304, for the release that you use, as follows:

- For R410, use Chapter 12 “Option 3—Install LIC and Recover Configuration.” Then use R410 Chapter 13 “Restoring the Operating System”
- For R370, use Chapter 15 “Option 3—Install LIC and Recover Configuration.” Then, use R370 Chapter 16 “Restoring the Operating System”
- For R360, use Chapter 10 “Option 3—Install LIC and Recover Configuration.” Then, use R360 Chapter 11 “Restoring the Operating System”

10. Restore user profiles.

11. Restore configuration.

12. Mount the tape saved from the CISC system in Step 1, and enter:

```
RSTLIB LIB(QUSRSYS QGPL)
```

13. To install QUSRSYS, QGPL, and the extended base support, enter:

```
GO LICPGM option 11
```

Use the IBM distribution media for this step.

14. Enter:

```
GO LICPGM
```

Select Option 1 to install Licensed Program Products from the DSLO tapes created in Step 2. On the Install Option panel, select option 3 for New Products.

15. Run: INZSYS

16. Run: RCLDLO *ALL

17. Use the tapes created in step 3 for these restore commands:

```
RSTLIB SAVLIB(*ALLUSR)
```

```
RSTDLO DLO(*ALL) DEV(tape-device-name)
```

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ(('/*')+
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))
```

18. Restore any user objects from QGPL and QUSRSYS that you need from the save in step 4. **Do not restore any IBM objects!** Use the RSTOBJ command and specify the objects by name. Specify those objects that were created during the test phase of the RISC system.

19. Run: RSTAUT USRPRF(*ALL)

20. Perform an IPL.

Note: Use this procedure only under the direction of an experienced system administrator and only as directed by IBM service personnel. If you need to better understand the procedure, contact the IBM AS/400 Support Line. Consult Line representatives are available to assist in reconciling libraries as an alternative to a system restore. Note that this option should only be considered in an “emergency” basis.

6.4 Journal Receivers and Previous Release Systems

For V4R2 systems to exchange journal receivers in a network involving previous release systems, apply PTFs to the earlier release system.

The following table lists the PTFs needed for the systems that remain at releases prior to V4R2.

Order This PTF	For System Remaining at Release
SF46569	V3R6M0
SF46570	V3R7M0
SF46571	V4R1M0

For more information, see APAR I110954. Refer to *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01.

Chapter 7. Licensed Program and PRPQ Backup and Recovery

As an AS/400 system administrator, your backup and recovery strategy will be unique to your environment. No two AS/400 systems are alike even in terms of what IBM software is installed. The software required on your AS/400 system is comprised at a minimum of the operating system, a set of licensed programs, and application code. Optional products installed include features of the operating system, as well as RPQ and PRPQ components.

This chapter discusses considerations involved with saving and restoring licensed program products (LPPs). Also included is a list of library names for many IBM licensed programs that may be on your system.

7.1 Licensed Program and PRPQ Considerations

The recommended method for saving licensed programs and PRPQs is to include these components as part of the full system backup strategy. This is done using the SAVE menu options.

Under most conditions, the SAVE menu options 21 and 22 are used to save licensed program product libraries and PRPQs for recovery purposes. In this case, using option 21 from the SAVE menu saves the licensed program product libraries as a part of the SAVLIB LIB(*NONSYS) command, and option 22 *System data only* from the SAVE menu as part of the SAVLIB LIB(*IBM) command. In addition, to ensure the correct saving of the portion of the licensed program product code stored in the integrated file system, run one of the following OBJ parameter strings for the SAV command:

```
OBJ(('/*) ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))
```

The previous command saves *all* integrated file system information.

```
OBJ(('QIBM/ProdData') ('QOpenSys/QIBM/ProdData')) UPDHST(*YES)
```

The previous SAV command string saves only the system licensed program product information. Again, the preferred method is to use option 21 and option 22 respectively from the SAVE menu.

Note: Two additional commands you need to run to ensure that *all* objects are saved for licensed program products are:

- SAVDLO—to save folder or document objects supplied with licensed program product code
- SAVSYS—because licensed program product commands can be in QSYS

Two reasons to save licensed programs separate from the operating system or a full system save are:

- To create separate tapes that allow for easier re-installation of a licensed program in the event the LPP becomes corrupted. The re-installation is easier because the PTFs for the program product are saved with the program code, and therefore, restored together. You do not have to apply PTFs separately after the LPP is restored.

- Distribution of program products to other AS/400 systems within a company's enterprise is easier to manage since fewer instructions and fewer tapes are involved.

In this case, to distribute licensed program products, use the SAVLICPGM command on one AS/400 system to save a Licensed Program with all PTFs applied and use the RSTLICPGM command to restore on another AS/400 system.

For more information on distributing licensed programs, refer to *Central Site Distribution*, SC41-5308.

7.2 Saving Licensed Programs

There are two reliable methods of saving licensed programs separate from the operating system on the AS/400 system:

- Using option 13 from the LICPGM menu
- Using the SAVLICPGM command

Note: Save files are used with SAVLICPGM but are not used with option 13 from the LICPGM menu.

You can access the Save Licensed Programs screen by entering the SAVLICPGM command and pressing the PF4 key. The example in Figure 25 shows how the SAVLICPGM command is used to save the DataPropagator/400 licensed program to a tape device named TAP01.

```

                                Save Licensed Program (SAVLICPGM)

Type choices, press Enter.

Product . . . . . LICPGM          5769DP2
Device . . . . . DEV              TAP01
                                + for more values
Optional part to be saved . . . OPTION *BASE
Release . . . . . RLS             *ONLY
Language for licensed program . LNG  *PRIMARY
Object type . . . . . OBJTYPE     *ALL
Volume identifier . . . . . VOL    *MOUNTED
                                + for more values
Sequence number . . . . . SEQ      *END
File expiration date . . . . . EXPDATE *PERM
End of tape option . . . . . ENDOPT  *REWIND
Save file . . . . . SAVF           _____
  Library . . . . .                *LIBL
Target release . . . . . TGTRLS     *CURRENT
Clear . . . . . CLEAR              *NONE
Data compression . . . . . DTACPR   *DEV

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 25. Using the SAVLICPGM Command to Save a Licensed Program

Note: It is important to note that when saving Client Access program products, be sure to also save the folders containing the program code. If you write your own CL programs for this backup, use the SAVDLO command for the associated folders. To determine the folder names associated with each program product, refer to *Backup and Recovery*, SC41-5304.

Using option 13 to *Save Licensed Programs* from the LICPGM menu ensures that the portion of the licensed product that resides in folders is saved. The LICPGM menu provides a list of licensed programs from which to work. Figure 26 shows the DataPropagator Relational for AS/400 licensed program selected for the save.

```

                                     Save Licensed Programs
                                     System:  SYSTEMXX

Type options, press Enter.
  1=Save

   Licensed Product
Option Program Option Description
-      5716DCT *BASE Language Dictionaries for AS/400
-      5716DCT 14   US Legal Dictionary
-      5716DCT 15   US Medical Dictionary
-      5716DCT 23   US English Dictionary
-      5769DFH *BASE CICS for AS/400          ment
-      5769DFH 1    CICS for AS/400 - Sample Applications
  1  5769DP1 *BASE DataPropagator Relational for AS/400
-      5769DS1 *BASE Business Graphics Utility for AS/400
-      5769FNT *BASE Advanced Function Printing Fonts for AS/400
-      5769FNT 1    AFP Fonts - Sonoran Serif
-      5769FNT 2    AFP Fonts - Sonoran Serif Headliner
-      5769FNT 3    AFP Fonts - Sonoran Sans Serif
-      5769FNT 4    AFP Fonts - Sonoran Sans Serif Headliner
                                     More...

F3=Exit  F11=Display status  F12=Cancel  F19=Display trademarks

(C) COPYRIGHT IBM CORP. 1980, 1998.

```

Figure 26. Using the Save Licensed Programs Menu to Save a Licensed Program

Note: SAVLICPGM does *not* save any of the user information used by the licensed programs.

7.3 Restoring Licensed Programs

There are two reliable methods of restoring licensed programs separate from the operating system on the AS/400 system:

- Using option 11 from the LICPGM menu
- Using the RSTLICPGM command

To access the Restore Licensed Program display, enter the RSTLICPGM command, and press the PF4 key. The example in Figure 27 on page 88 shows how the RSTLICPGM command is used to restore the DataPropagator/400 licensed program from a device named TAP01.

```

                                Restore Licensed Program (RSTLICPGM)

Type choices, press Enter.

Product . . . . . LICPGM          5769DP2
Device . . . . . DEV              TAP01
                                + for more options
Optional part to be restored . . OPTION  *BASE
Type of object to be restored . . RSTOBJ *ALL
Language for licensed program . . LNG     *PRIMARY
Output . . . . . OUTPUT          *NONE
Release . . . . . RLS            *FIRST
Replace release . . . . . REPLACERLS *ONLY
Volume identifier . . . . . VOL      *MOUNTED
                                + for more values

Sequence number . . . . . SEQNBR      *SEARCH
End of tape option . . . . . ENDOPT   *REWIND
Save file . . . . . SAVF            _____
Library . . . . . _____         *LIBL

                                                More...

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 27. Using the RSTLICPGM Command to Restore a Licensed Program

Using option 21 from the LICPGM menu ensures that the restore requests the portion of the licensed program that resides in folders. The LICPGM menu provides a list of licensed programs from which you can work.

The RSTLICPGM command is also used to install a secondary language for a licensed program, as is option 21 from the Work with Licensed Programs (LICPGM) menu.

7.4 Special Considerations to Install and Restore Licensed Programs

Some licensed program products or PRPQs have special installation instructions. Refer to the information provided with the installation media for unique product requirements when restoring the program on the target system. Three considerations are described in this section.

7.4.1 LICPGM Menu Does Not Manage all Licensed Programs

You cannot assume that the AS/400 installation menus can be used for managing all licensed program installations. For example, Lotus Domino for AS/400 is installed with the LODRUN utility from the Lotus CD. It is not installed using the LICPGM menu. Therefore the LICPGM menu cannot be used to display, install, delete or save the Lotus Domino for AS/400 product. The Lotus Domino for AS/400 product “shows up” when using the DSPSFWRSC command (with a Resouce ID of 5769-LNT).

7.4.2 User Profile Authority with Licensed Programs

The user class parameter value on the user profile needs *USER authority or higher for a user to work with the LICPGM menu. This is true *even if* the user profile has *ALLOBJ authority. A user with a user class of *PGMR but no special authorities *can* work with licensed program options. If the user profile does not have higher than *USER authority, a blank screen (a screen with no options) appears on the LICPGM menu.

Recommendation

Ensure the user profile using the LICPGM menu has *USER (user class) authority.

In other words, the user class parameter overrides any special authorities depicted on the user profile.

7.4.3 Multi-National Considerations with Directory Names

The SAVE and RESTORE menu options 22 and 23 run SAV and RST commands which process path names with lower case characters. For multinational systems or systems running with CCSID 290 or 5026, it is important to note that **at the time of this writing, path names with lower case letters** will cause options 22 and 23 to fail with an *object not found* message. This means that the SAV is not processed correctly when the job CCSID is 290 or 5026, for directories named as:

```
OBJ('/QIBM/ProdData') (/QOpenSys/QIBM/ProdData)
```

These CCSIDs are most commonly used in Japan.

As a circumvention until this problem is resolved in a future release, the 5035 CCSID can be used to map lower case English characters to Japanese characters. To change the CCSID of the job to 5035 before running options 22 or 23 from the SAVE or RESTORE menus, use the following command:

```
CHGJOB CCSID(5035)
```

This command also works when entering the SAV or RST command on the command line if directory names contain lower case letters.

7.5 Licensed Program Library Names for IBM Supplied Libraries

The following table identifies the library names included with the licensed programs compatible with V4R2. This table is useful to understand what libraries and folders should not be deleted or which should be backed up.

Table 7 (Page 2 of 3). Licensed Programs (Including Library Names) Compatible with V4R2

Product Number	Library Name	Product Name
5769TC1	QTCP	TCP/IP Connectivity Utilities for AS/400
5716UB1	QUMCONFER, QUMBCEDS, QUMBCOS2, QUMPPPOS2, QUMPPWIN	Ultimedia Business Conferencing for AS/400
5716US1	QUMHCATL, QUMPIOSS, QUMPISAM, QUMBLOSS, QUMBLSAM, QUMHCATL	Client Access Ultimedia Tools for AS/400
5716VG1	QVGEN	VisualGen Host Services for AS/400
5769WP1	QOFC, QBBCSRCH, QDOC (folder)	OfficeVision for AS/400
5763XB1	QIWSP, QIWSPS, QIFWSPD, QRUMBA, QRUMBAD, QUMSFWIN, QIWSP, QIWSPS, QIWSPD, QDOC (folder)	Client Access for DOS with
5763XC1	QPWXCLIB, QPWXCNN, QPWXCWND, QPWXCRB, QPWXCRBD, QPWXCPC, QPWXCGY, QPWXCUM, QPWXCGA, QPWXC50	Client Access for Win
5763XD1	QWIN32	Client Access for Windows 95/NT
5763XF1	QIWS2, QIWS2S, QIWS2D, QRUMBA2, QRUMBA2D, QCM400, QGYOS2, QUMSFOS2, QIWS2, QIWS2S	Client Access for OS/2/RUMBA
5763XG1	QPWXFOS2, QPWXGRB, QPWXGPC, QPWXGGY, QPWXGUM, QPWXGGA	Client Access Optimized for OS/2
5763XK1	QWIN16, QWIN16S, QWIN16D, QPC5250K, QPC52S0T, QPC5250P	Client Access Enhanced for Windows 3.1
5763XL1	WSF, QIWSFS, QIWSFD	Client Access for DOS
5769XW1	QCA400N	AS/400 Client Access Family for Windows
5769XY1	QCA400Y	AS/400 Client Access Family
5769XZ1	QXZ1	OS/2 Warp Server for AS/400

Table 7 (Page 3 of 3). Licensed Programs (Including Library Names) Compatible with V4R2

Product Number	Library Name	Product Name
5769SS1	QSYS, QGPL, QUSRSYS, QSYS2, QQALIB, QHLPYSYS, QSYSDIR, QMGU, QSSP, #DFULIB, #SEULIB, #DSULIB, #SDALIB, #CGULIB, QSYS38, QUSRTOOL, QFNTCPL, QSYSV3R7M0, QSYSV3R6M0, QSYSV3R2M0, QSYSV3R1M0, QMU400, QIWS, QSYSINC, QGDDM, QCPA, QUMEDIA, QAFPLIB, QMSE, QM36, QSYSLOCALE, QSR, QSOC, QFSNOTES, QFPNTWE, QSMP, QDB2MS, QFLOWMARK, QSYCGI, QGY	Operating System/400

Note: The *Central Site Distribution*, SC41-5308, guide also contains a list of library names for IBM products.

7.6 Restoring Commands from Licensed Program Libraries to QSYS

To better administer the AS/400, it is important to know what libraries are associated with each IBM product. Sometimes there are multiple libraries involved for any given product. As you can infer from Section 7.5, "Licensed Program Library Names for IBM Supplied Libraries" on page 89, this makes it easier to lose track of how products are packaged. For example, *CMD objects stored in QSYS can be lost when recovering a system from distribution media. If this occurs, use this section to restore the licensed program command objects to QSYS.

To view the library names associated with a given IBM product, press F11 from the DSPSFWRSC command output. This information helps to find the appropriate library to restore the licensed program commands.

To list and restore commands to QSYS for a given licensed program, complete these steps:

1. On the AS/400 command line, type:


```
DSPSFWRSC
```

Press ENTER.
2. When the list appears, press F11. This brings up the library names for the licensed programs.

3. For *all* products listed, *except* 57nnSS1, write down the names of the product libraries you want to restore commands from. For example, if product 57nnSV1 SystemView for AS/400 is listed, pressing F11 shows that the related libraries are QSVBASE and QSVLNCH. Do not write down library names for 57nnSS1 products.

4. For each library name that is recorded in step 3, issue this command:

```
WRKOBJ OBJ(xxxxxxx/*ALL) OBJTYPE(*CMD)
```

The variable xxxxxxxx is the name of the library. This step lists all the commands associated with that product.

5. To copy these commands to QSYS, type a 3 beside each command listed. Then, type the following parameter on the command line:

```
TOLIB(QSYS)
```

Press ENTER. This copies each command selected with option 3 from the product library to the QSYS system library.

6. Repeat steps 4 and 5 for each library recorded in step 3.

Attention

Do not copy anything from QUSRSYS.

7.7 For More Information

For more information about saving and restoring licensed programs, refer to *Software Installation*, SC41-5120, and *Backup and Recovery*, SC41-5304. A list of library, folder and media files can be found in the *Software Installation*, SC41-5120, publication as well.

Chapter 8. Save, Restore, and System Performance for Availability

Apart from availability considerations described throughout this redbook (specifically Chapter 10, “Work Management for System Availability” on page 139), there are a number of performance factors on which the system administrator has influence to improve system availability. These factors include performance tuning parameters, the choice of hardware, and configuration options for performing save and restore processes.

The purpose of this chapter is to:

1. Provide performance information on selected IBM tape drives to help you plan your save and restore configuration.
2. Describe some system tuning areas you can control to give the best overall system performance.

8.1 Save and Restore Performance

To plan for implementing appropriate tape drives for your availability and recovery requirements, you need to understand the tape hardware and capabilities of that hardware. Different tape drives and IOPs have different capabilities and performance ratings. Note that a slower rated tape drive is actually faster for some environments.

Many factors influence the observable performance differences of save and restore operations. These factors include:

- Using data compaction on tape
- Using Use Optimum Block Size (USEOPTBLK)
- Using data compression
- Tape drive options
- The tape input/output processor (IOP)
- Placement of the tape IOP in the system
- Type of workload (size of files, mix of user workloads)

Table 8 on page 96 shows the tape drives that were tested in the IBM laboratory and are referenced later in this section and in Appendix C, “Save and Restore Rates of IBM Tape Drives for Sample Workloads” on page 379. This table shows the rates for each drive.

8.1.1 Data Compaction for Tape

Some AS/400 IOPs provide data compression in the hardware data path, which is known as hardware data compression (HDC). With HDC, the data rate and capacity of the attached tape device can be enhanced.

Data compaction is only available at the hardware level. Therefore, if you want to use it, the tape drive you choose needs to support it.

Tape drives that support data compaction, as shown in Table 8 on page 96 include:

- 6390
- 7208-342
- 6385

- 3570
- 3490
- 3590

Note that the different rates are obtained by different combinations of drives and parameters on the save and restore commands.

IBM conducted a study about the affect of data compaction on save and restore using 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 200MB and 2GB workloads is based on this ratio for tape drives that use LZ1 for this compaction algorithm. The same data compacts at about 1.8 to 1 for IDRC drives.

Table 8. Tape Drive Ratings

Drive	Tape Drive Rate (MB/S)	DATA COMPACTION (COMPACT)	Optimum Blocking (USEOPTBLK)
6380	0.3	0.0 compaction not available	
6382/6383	0.4	0.0 compaction not available	
6390	0.5	1.8	
7208-342	3.0	1.8	YES
6385	1.5	1.8	
3570	2.2	2.5	YES
3490	3.0	1.8	
3590	9.0	1.6 ¹	YES
1 The 3590 does not make use of full compaction. The figures listed are simulated for the 3590.			

Note: For interchange and compatibility with IOPs that do not provide hardware data compression, the HDC algorithm is implemented at the software level with software data compression (SDC). SDC increases performance for slower tape devices. For the highest speed tape drives, however, SDC severely limits performance.

HDC and SDC are controlled by the data compression (DTACPR) parameter of the save commands:

- DTACPR(*DEV) (the default) activates HDC if it is supported, otherwise HDC is not used.
- DTACPR(*YES) activates HDC if it is supported, otherwise it uses SDC.
- DTACPR(*NO) does not use either HDC or SDC.

Recommendation

Customers that have 3490 Cnn tape drives attached with a 2644 IOP should specify DTACPR(*YES) for maximum performance.

If you upgrade to a 3590 or a 3490 Enn tape drive with a 6501 IOP, change the data compaction parameter to the default of DTACPR(*DEV). If DTACPR is specified as *YES for drives attached to the the 6501, software data compression (SDC) is used and performance is not as efficient.

8.1.2 Use Optimum Block Size

The use optimum block size (USEOPTBLK) parameter sends a larger block of data to tape drives that can take advantage of the larger block size.

Tape drives that support optimum blocking, as shown in Table 8 on page 96, include:

- 7208-342
- 3570
- 3590

Each block of data that is sent has a certain amount of overhead associated with it. This overhead includes:

- Block transfer time
- IOP overhead
- Drive overhead

The block size does not change the overhead of the IOP or drive. The number of blocks does affect the overhead of the IOP or drive. For example, sending eight small blocks results in eight times as much IOP and drive overhead. If the same amount of data is sent in one large block, the result is one times the overhead.

With the larger block size, the overhead of the IOP and drive become less significant. This allows the actual transfer time of the data to be the limiting factor. In our example, eight software operations with eight hardware operations are essentially a software and hardware operation when USEOPTBLK(*YES) is specified. The result is usually a significantly lower use of the CPU, which also allows the tape device to perform more efficiently.

8.1.3 Data Compression

Data compression (DTACPR) is the ability to compress strings of identical characters and mark the beginning of the compressed string with a control byte. The control byte serves as a “counter” of replicated characters. Strings of blanks from 2 to 63 bytes compress to a single byte. Strings of identical characters between 3 and 63 bytes compress to two bytes. If a string cannot compress, a control character is added that expands the data. DTACPR is usually used to conserve storage media.

If the IOP does not support data compression, software performs the compression, which can require a considerable amount of processing power. One of the exceptions to this is the 2644 IOP card for the 3490 tape drive. Here, the IOP is designed to use compression as noted in Section 8.1.4, “Save and Restore Tips and Techniques for Better Performance” on page 98 of this chapter.

The purpose of hardware data compression (HDC) is to increase the data rate and capacity of the attached tape device. For interchange and compatibility with IOPs that do not provide HDC, the HDC algorithm is implemented in the system SDC. On CISC systems, the I/O bus rate is 8MB for tape operations. The maximum 3490 rate is 4.9MB. The maximum 3590 rate is 5.2MB. Therefore, if the 3490 or 3590 are placed on a bus with a significant portion of the DASD (or with each other), performance is adversely affected.

For RISC systems, the I/O bus rate is 16MB or 24MB for tape operations. The maximum 3490 rate is 4.9MB. The maximum 3590 rate is 9.0MB. On RISC

systems with the improved I/O bus rate, DASD and high performance tape can be mixed on an I/O bus as long as they do not exceed the I/O bus bandwidth. The following recommendations explain how to maximize save performance to tape.

- The tape subsystem bandwidth is equally important as is the I/O bandwidth. If the device has a controller that attaches multiple host systems and multiple devices, they compete for device bandwidth in the same manner. If two 3490 devices are connected to one controller (3490 C2A), the controller must share the data path between two devices. This cuts throughput to less than half.
- For optimum performance when running in a restricted state (saves, restores, RCLSTG, and so forth), vary off all communication lines to prevent unnecessary system processing of retries or reconnections. Take the time to vary off all configuration objects during the time of the save or restore operation. Otherwise you potentially suffer degraded performance when in a restricted state.
- Use concurrent saves to reduce the time to save. Some overlap processing can be done during saves or restores. For example, use two different tape drives at the same time for SAVLIB LIB(*ALLUSR) and SAVDLO. Or, use SAVLIB on two different lists of libraries with two different tape drives at the same time.
- If you have many DLOs, consider putting DLOs in different ASPs. This not only increases the total number of DLOs that you can have on your system, but also improves recovery time in the event that a restore is required.

Refer to Chapter 5, "Save and Restore for Availability and Recovery" on page 49, and Section 5.12, "Product Preview for Save Restore Enhancements" on page 73, for considerations when planning data storage solutions.

8.1.4 Save and Restore Tips and Techniques for Better Performance

The following recommendations improve the performance of save and restore operations:

1. To achieve maximum tape performance, the system must have a balance in the placement of high performance loads on the system. Since the data written to tape must come from DASD, the DASD and tape operations must be equal or less than the system bus bandwidth. Two high-performance tape drives on the same bus can have the same affect. During a save-while-active function, large LAN or workstation networks can compete with the tape for system bus bandwidth. Making use of the latest DASD architectures with the fastest service times is an advantage for save and restore performance improvements. (Service time is the time the read and write heads take to access a given piece of data).
2. The choice of using hardware data compression (DTACPR), data compaction (COMPACT), and the USEOPTBLK parameter are important to get the best performance from your tape drives. The settings in the following table work best on average data.

Table 9 (Page 1 of 2). Recommended Settings for Optimum Performance

Tape Device	DTACPR	COMPACT	USEOPTBLK
6380	*YES	*NO	*NO

<i>Table 9 (Page 2 of 2). Recommended Settings for Optimum Performance</i>			
Tape Device	DTACPR	COMPACT	USEOPTBLK
6382/6383	*YES	*NO	*NO
6390	*DEV	*DEV	*NO
7208-342	*DEV	*DEV	*YES
6385 ¹	*DEV	*DEV	*NO
3570	*DEV	*DEV	*YES
3490 (2466 IOP)	*YES	*NO	*NO
3490	*DEV	*DEV	*NO
3590	*DEV	*DEV	*YES
¹ The 6385 can experience performance degradation if the system cannot keep its buffer filled. In this case, you may find better performance with DTACPR(*NO) COMPACT(*NO).			

Experiment with these parameter settings to get the best performance for your environment.

3. Using USEOPTBLK significantly improves performance on later model tape drives. This is especially true where the system's CPU is subject to a heavy workload. The newer drives are designed to take advantage of the USEOPTBLK parameter. Remember that for V4R2 systems, USEOPTBLK is the default. On systems prior to V4R2, you may want to consider specifying USEOPTBLK(*YES) or change the command default (as is described in Section 5.6, "Use Optimum Blocking for Save and Restore" on page 57).
4. Adjusting the system configuration is most beneficial for fast tape drives. For the fastest save times, the tape IOP may not be installed where it can be used as an alternate IPL device. This is a consideration on systems prior to V4R1. In that case, you need a slot in the main system unit to move to if you ever need to restore your system. On V4R1 or later systems, the IPL device is no longer restricted to the main system unit. Refer to Chapter 3, "Availability Options Provided by Hardware" on page 23, in this redbook for more information about optimizing the tape drive placement.
5. Write capability indicates the rate at which the system can send data to a tape device across the bus. For bus speed, the primary concern is with write capability. On V4R1 or later systems, the 3590 is the only tape drive that can outperform some of the buses. However, the concepts in this section can help other tape drives perform better by avoiding data collisions. For most of the buses, the system can read data from devices much faster than they can write to them.
6. The placement of the tape and DASD IOPs are important to save and restore performance. There are several factors to consider when placing the IOP card:
 - The bus in the main system unit varies in speed depending on the system model you have.

To avoid collisions between data that is read or written to disk and data that is read or written to tape, spread DASD and tape utilization across multiple IOPs and buses. For example, IBM testing has determined that 40 or more DASD arms spread across three buses works well for a single tape drive.

- To drive multiple tape drives, create balanced user Auxiliary Storage Pools (ASPs). You must know which tape drive is being fed by which ASP. This ensures that the tape IOP is placed on the proper bus to avoid data flowing from the other ASPs. A good candidate is a backup or restore of several large data files. IBM Consult Line services are available to help determine the best configuration for your environment.
- Optical bus cables come in different lengths. The length of optical bus cables can affect save and restore throughput.
- When you design your files, keep save performance in mind. One file with 10 000 members is faster to save and restore than 10 000 files with one member each.

Refer to Appendix C, “Save and Restore Rates of IBM Tape Drives for Sample Workloads” on page 379, for save and restore measurements for selected IBM tape drives.

8.2 Defective Device or Media Considerations

If hardware components or media cause a high number of errors, performance degrades. To thoroughly analyze a performance problem, examine device or media recovery actions using the error logs and media statistics. Review the logs proactively to catch errors before they cause a hard failure.

8.3 What a User Can Do to Influence System Performance

More work is done and a longer window is available for planned outages when a system performs at its best. Good performance on an AS/400 system is, at first glance, not considered an availability function. However, unavailability is when the system or your applications negatively affect workload throughput and the efficiency of its operations and management.

Users have influence in many areas of system performance. Those described in this section include:

- Automatic Performance Adjustments
- Altering Shared Pool and Priority Values in Conjunction with the Automatic Tuner
- Tuning the System Tuner
- Tuning Parameters for Batch
- Dynamic Priority and Controlling CPU Intensive Jobs

8.3.1 Automatic Performance Adjustments—It Is Worth Another Look

The QPFRADJ system job takes a snapshot of main storage every 20 seconds. QPFRADJ then compares the last snapshot with the running average calculated from data acquired during the past 60-second time span. From this comparison, it determines whether to add or subtract memory in shared pools. QPFRADJ also adjusts the activity levels within shared pools. As is the key to good memory usage, QPFRADJ tries to keep the total faulting rates in all shared pools as low as possible.

Due to the large range of AS/400 processors and an increasing variance in the complexity of user applications, paging guidelines for user pools are not

published for V4R2 systems. Only machine-pool and system-wide guidelines (the sum of faults in all the pools) are published.

The written guideline that the QPFRADJ system job uses is the page fault guideline for the ***MACHINE POOL**. For user-defined memory pools, the guidelines are calculated dynamically. All processor models use the same adjusting algorithm for pools and activity levels. Entry-level systems are treated the same as high-end systems.

The algorithm *is* sensitive to the server models, since server models are designed primarily for batch throughput. By default, the ***INTERACT pool** has a lower adjustment priority on server models.

Recommendation

Set QPFRADJ to a value of 2 (adjust at IPL and during run time) or 3 (adjust at run time only) if you are not prepared to manually set pool sizes and monitor page fault rates.

Note: The automatic tuner process has undergone many algorithm changes since it first appeared in Version 2 of the operating system. Performance adjustments are made more quickly to changes within each pool than in releases prior to V3R1. Do not let your previous experience with the CISC release implementations of automatic adjustment cause you to set system value QPFRADJ to 0 (off) without reconsidering activating the system tuner. QPFRADJ can help manage your shared storage pool sizes and activity levels.

For more information, see the Performance Tuning chapter of the *Work Management Guide*, SC41-5306.

8.3.2 Altering Shared Pool and Priority Values with the Automatic Tuner

The Work with Shared Pools (WRKSHRPOOL) and the Change Shared Pool (CHGSHRPOOL) displays show the values that the QPFRADJ automatic system tuning job uses to assign storage and priority to shared pools. You can change the values from these same displays. After selecting F11 to *Display pool data*, you can choose to alter any of these values:

- Pool priority value
- Pool minimum and maximum storage size specified as a percentage of the total main storage
- Pool page faults-per-second range for each shared pool and jobs running in that pool

Figure 28 on page 102, the Work with Shared Pools display, shows where the pool priority, minimum and maximum storage size, and faults-per-second range values are located.

```

                                Work with Shared Pools
                                System:  SYSTEMXX

Main storage size (K) . :      393216

Type changes (if allowed), press Enter.

-----Size %-----  -----Faults/Second-----
Pool      Priority  Minimum  Maximum  Minimum  Job  Maximum
*MACHINE  _____  12.92   100      10.00   .00  10.00
*BASE     _____  .50     100      10.00   2.00  100
*INTERACT _____  10.00   100      5.00    .50   200
*SPOOL    _____  1.00    100      5.00    1.00  100
*SHRPOOL1 _____  1.00    100      10.00   2.00  100
*SHRPOOL2 _____  1.00    100      10.00   2.00  100
*SHRPOOL3 _____  1.00    100      10.00   2.00  100
*SHRPOOL4 _____  1.00    100      10.00   2.00  100
*SHRPOOL5 _____  1.00    100      10.00   2.00  100
*SHRPOOL6 _____  1.00    100      10.00   2.00  100
More...

Command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F11=Display text
F12=Cancel

```

Figure 28. An Example of Work With Shared Pools Panel

Priority Parameter: The priority of a shared storage pool is similar to the priorities of other shared storage pools. The priority ranges from 1 to 14, where 1 is the best priority and 14 is the worst priority. The system uses this value if the QPFRADJ system value is set to 2 or 3. When you enter a value of *DFT, the default value appears.

On V4R2 systems:

- The *MACHINE pool has the best priority value (1) compared to other shared storage pools.
- On server models, *INTERACT pool is shipped with priority 2 and *BASE pool with priority 1.
- On non-server system models, *INTERACT pool is shipped with priority 1 and *BASE pool with priority 2.
- All other shared storage pools are shipped with priority 2.

Size Percentage Parameter: This parameter indicates the percentage of the total main storage that is available to allocate to this storage pool. The system uses the size parameter value if the QPFRADJ system value is set to 2 or 3. Enter a value of *DFT, and the default value appears. The default values are highlighted so determining any change to the setting is easy—the changed value is highlighted on the user’s display.

Minimum This is the minimum amount of storage to allocate to this storage pool as a percentage of the total main storage.

On most systems, the machine pool has the highest default value for “minimum percentage of total storage.” For the machine pool, this value changes throughout the day if the QPFRADJ job is active. QPFRADJ takes into account the minimum amount of storage that must be reserved for system use. The WRKSYSSTS command output shows the amount of storage.

On AS/400 system models (not server models), *INTERACT pool has the largest amount of minimum storage aside from the machine pool. *INTERACT defaults to 10% of the available main storage. All other shared storage pools have a minimum of 1% of the available main storage.

The *BASE storage pool is calculated using the system value QBASPOOL. For the WRKSHRPOOL example shown in Figure 28 on page 102 with a system value of QBASPOOL(2000) and total main storage of 393216K, the minimum size for *BASE is calculated as follows:

$$(2000/393216) \times 100\% = .50\%$$

This is truncated to two decimal positions.

Recommendation

Size the *BASE pool to at least 5% of the total installed main storage. Change the default on the *WRKSHRPOOL* screen to 5% for *BASE.

Maximum This value indicates the maximum amount of storage to allocate to this storage pool as a percentage of the total main storage. The actual maximum amount of storage that is assigned is determined by this percentage and affected by the amount of storage allocated to other active pools.

Faults-per-Second Parameter: This parameter indicates the number of page faults per second for the system to use as a guideline for this pool. The system uses this value if the QPFRADJ system value is set to 2 or 3 for automatic performance adjustment. Enter a value of *DFT to view the default value. Again, all the default values are highlighted on the screen.

Minimum Page faults per second to use as a guideline for this storage pool.

Job Page faults per second for each “active” job to use as a guideline for this storage pool. In this calculation, a job is counted as active if it used any CPU cycles in the last twenty seconds.

Maximum Page faults per second to use as a guideline for this storage pool.

For *INTERACT with 20 active jobs, the paging guideline calculated by the system is:

$$5 + 20 \times .50 = 15 \text{ faults/second}$$

QPFRADJ does not adjust its maximum page fault guidelines based on processor speed groups. This processor-dependent set of “good” values is documented in the *Work Management Guide*, SC41-5306, for each pool. QPFRADJ defaults to maximum per pool values of 100 for all pools except *MACHINE pool and 200 for *INTERACT. The high-end models can tolerate higher page fault rates before page faults per second impact performance. Therefore, you may want to increase the *Maximum faults/second* value on the WRKSHRPOOL panel.

Remember, a pool is in need of more storage if its page faults per second are beyond a threshold value. Also, a pool needing more storage cannot receive additional storage unless another storage pool can release it without reaching its own threshold values of page faults per second and minimum size.

8.3.3 Tuning the System Tuner

The following items represent actions that you can perform to fine tune the system performance adjuster (QPFRADJ):

- Increase the priority of selected pools. If you have a pool for “mission critical” jobs, adjust the pool’s priority relative to other pools. To do this, lower the adjusting priority of all other pools. Make sure that you allow the machine pool to remain top priority. You may also decrease the faults-per-second guideline to ensure that the pool has more memory per job running in that pool.
- Decrease the paging guideline of selected pools.
- Set a minimum or maximum size for selected pools. By setting a minimum size of a pool, you indicate the lowest amount of main storage available for the pool. The same applies to the maximum pool size.
- Remove a shared pool from QPFRADJ adjustment considerations. You can force the QPFRADJ system job to leave a pool alone, despite what it may do otherwise. This is done by setting the *minimum* and *maximum* sizes to equal values.

This action is useful to keep a pool from being reduced too much when jobs are idle (such as, over a lunch break). Setting a minimum size ensures there is enough storage for users when they become active on the system again.

- Switch back to fixed guidelines. You can return to the “old,” fixed guidelines for any or all of the pools by setting the amount of faults to 30 faults per second, for example, regardless of how many jobs are running in the pool. To use the old guideline of 30, set *Minimum* to 30, *Job* to 0, and *Maximum* to 30.

8.3.4 Tuning Parameters for Batch

Expert Cache is an OS/400 function that attempts to anticipate a need for data in each pool so that applications do not have to wait for disk access to reach the data. Expert cache is turned on for selected pools by setting the shared pools **Paging Option** to ***CALC**. Specific batch environments that are DASD I/O intensive and process data sequentially realize significant performance gains by using expert cache. An example of this is a save and restore job.

Expert cache does not provide improvement unless the following statements are true:

- The application running is disk intensive, and disk I/Os are the limiting factor in throughput.
- The system has sufficient main storage to keep the page faults within guidelines.
- The processor is under-used (usage is at less than 60%).

Spare CPU cycles result from:

- Expert cache reducing the disk I/O, therefore, reducing disk access time, and providing more information for the CPU to process
- Accessing data in main storage faster than on disk

Caching in main storage reduces the average disk access time. This process allows the CPU to process more work.

8.3.5 Dynamic Priority and Controlling CPU Intensive Jobs

Dynamic priority scheduling involves a system value that allows you to turn on and off the dynamic priority scheduler. The system task scheduler uses this value to determine the algorithm for scheduling jobs running on the system. Dynamic priority scheduling ensures that a CPU-intensive job does not saturate the CPU resource. This feature is especially beneficial when there are many interactive jobs running on the system.

Dynamic priority scheduling ensures that batch jobs acquire CPU resources without significantly impacting interactive jobs. Any jobs that are given a priority of 0 through 9 are placed on a high priority flat curve. The task dispatcher checks this “priority class” first before checking any other dynamic priority classes. A looping job running with a priority of 0 through 9 causes serious performance problems to all users on the system.

One instance that can cause intensive CPU activity is when a user interactively performs a *wild card* search (for example, looking for the occurrence of the character string DMT*). Control the use of such searches and place proper authority on the commands or processes that are involved. However, the performance impact is less when dynamic priority scheduling is active.

Recommendation

Set the Dynamic Priority Scheduling system value QDYNPTYSCD to 1, which is the system default). This allows for optimum distribution of resources and overall system performance.

8.3.5.1 Freeing a Hung System when the Console is Not Available

There may be an occasion when the CPU is so overworked that the entire system slows down and even the high-priority console does not respond to operator requests. The “hog hunter” function available through the control panel of the processor (the operator panel) can help you uncover user jobs that are *hogging* the CPU.

The hog hunter function invokes a vertical licensed internal code (VLIC) process that looks for a high priority, CPU blocking job. If it detects a *user* job, it lowers the priority level immediately by 20 points. For example, if the priority level is 15, hog hunter changes it to 35 upon finding it. Hog hunter does not alter a system job’s priority.

When all workstations (including the console) are input inhibited and the System Activity light is solid, the following steps relieve CPU activity if the congestion is caused by a user job.

1. Select manual mode.
 - On a white (CISC) system, turn the key to manual.
 - If you are on a black Model 200 system, select function 02 on the front panel.
 - If you are on a black Model 300/500 system, press the mode button.
2. Select function 25 and press the Enter key.
3. Select function 26 and press the Enter key.

4. Select function 76.

One of six SRC codes in the control panel lights appears:

1. D6000652 or A9002052 service reference code (SRC) indicating a job matching the criteria for a CPU hogging job is found and its priority is lowered by 20 points.

If the workstations are still input inhibited, select and enter function 76 again until you free a workstation or no longer receive one of the SRC codes.

2. D6000653 or A9002053 indicating that the job detected, which is hogging the system, is a system job such as QSYSARB or QLUS. Although the priority of a system job is not altered, the user now knows what to investigate when the system performs IPL or frees up again.

Multiple uses of function 76 continue to lower the priority of any user job matching the criteria by 20 points to a maximum of 99.

3. D6000654 or A9002054 indicate there are no jobs that meet the hog hunter criteria and, therefore, no changes are made.

Note: If no user job is encountered, request a main storage dump. The dump can be analyzed to determine where the hang is and perhaps what caused it.

If a priority reduction is successful, message CPD0940 is sent to QHST and QSYSOPR or QSYSMSG message queues indicating:

```
Job &3/&2/&1 using CPU resources is changed . . .  
The priority is changed from &4  
to &5 because option 76 is entered on the control panel
```

Note: To clear the SRC, sign on the console.

A VLIC log (or VLOG) is also produced. Depending on what type of job the hog hunter encounters, you can find VLOG entries for:

- VL17000301 Resource Management
- VL17000302 Resource Management
- VL17000303 Resource Management

Consult Line personnel can interpret the VLOG entry to determine what the CPU intensive job is.

Note: The hog hunter function is included in the system code on R320 systems and later. For R310 systems, apply these PTFs to install the hog hunter code:

- MF11581
- MF11982
- MF11960
- SF29221

Refer to *AS/400 Licensed Internal Code Diagnostic Aids—Volume 1*, LY44-5900, for information on the use of function 76, and *AS/400 Licensed Internal Code Diagnostic Aids—Volume 2*, LY44-5901, for a description of VLOGs.

8.4 For More Information

A comprehensive list of tuning tips information can be found in the redbook *AS/400 Performance Management V3R6/V3R7*, SG24-4735, and *Performance Tools/400*, SC41-5340, as well as the “Performance Tuning” chapter of the *Work Management*, SC412-5306, manual.

Chapter 9. Tools for Automating System Management Functions

In this chapter, we describe the main functions and features of various tools which help the system administrator in their multifaceted job of managing change, problems, operations, and data. There are tools and procedures to make it easier to recover from a disaster, automate message handling, and ensure that jobs run according to plan.

The following tools are described in this chapter:

- ADSTAR Distributed Storage Manager/400 (ADSM/400)
- Backup Recovery and Media Services for AS/400 (BRMS)
- WRKASP Utility
- Print System Information Tool
- OS/400 Job Scheduler (part of the operating system)
- IBM Job Scheduler for OS/400
- IBM SystemView System Manager for AS/400
- IBM SystemView Managed System Services for AS/400
- Automating Message Management
- OS/400 Alert Support
- Automating Security Management

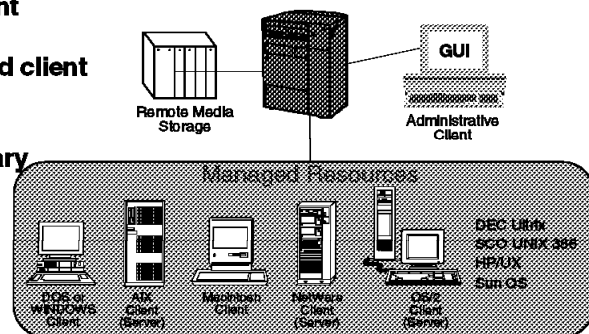
9.1 ADSTAR Distributed Storage Manager/400 (ADSM/400)

Backing up a workstation can be a tedious process requiring constant attention. ADSM helps those administrators responsible for backing up a large number of workstations.

ADSTAR Distributed Storage Manager/400

Data Management

- **Enterprise backup and recovery**
-- AS/400 Server and attached clients
- **Tape library management**
- **Automated, policy-based client backup and recovery**
- **Archival and media library management**



Protection of Mission-Critical Data

Figure 29. ADSM/400

Because workstations are “personal,” they go against the formal nature that a strict backup process requires, placing the responsibility for managing the workstation onto the shoulders of the end user. Typically, users are ineffective at saving critical data and, in the worst cases, never back up their workstations at all.

ADSM is designed specifically for client/server save and restore operations to perform backups of client data to a server system (in this case, the AS/400 system). ADSM is a cross-platform product supporting over 30 types of clients, including DOS, Microsoft Windows, and OS/2 systems.

9.1.1 ADSM/400 Customer Scenario

In a large AS/400 installation, an administrator is in charge of defining the setup for backing up all the users’ PCs to the AS/400 server with ADSM. This backup is set up to run automatically every Friday from each client. The backup contains whatever files are specified in an ADSM Options File for each client.

One Tuesday afternoon, the company’s marketing manager calls the administrator because he suddenly discovers that he has lost the company’s marketing strategy Freelance presentation from his PC. He is scheduled to present the strategy to the CEO on the following day. What he wants to know is whether an older version of his presentation has been saved on the AS/400 server, so that he can make a few last-minute corrections, or whether he has to start all over again.

The administrator performs the following tasks:

- Checks that the marketing manager’s PC is in fact an ADSM client.
- Uses ADSM facilities to check when this PC was last backed up.
- Helps the marketing manager restore the damaged file back to his PC.
- Shows the marketing manager how to do an adhoc backup of the presentation to the AS/400 server after making corrections.

When a colleague of the marketing manager hears about this happy ending, he decides that he also wants his PC backed up every Friday. So the administrator adds this new user to the backup schedule. As the use of ADSM grows to protect PC users’ data, the administrator’s responsibility increases. He defines another user as an administrator so that there are two to share the workload of helping users with their backups.

9.1.2 ADSM Elements

ADSM/400 consists of three elements:

The backup clients The backup clients are user PCs or PC servers that have data that needs to be backed up.

The administrative client The administrative client is a user who controls and monitors the whole backup/recovery process for the PC users and manages the server.

The server The server is where the backed up data is stored.

9.1.2.1 Backup Clients

ADSM/400 supports backup of a large variety of IBM and non-IBM workstations. In addition, ADSM/400 provides a backup agent for the Lotus Notes database running on the Integrated PC Server (IPCS).

The following are the main functions that ADSM/400 offers to client workstations:

- Backup function

The backup function is a version-based save that is performed in one of three ways:

1. Incremental (such as SAVCHGOBJ on the AS/400 system)
2. Selective (nominating a file to save based on file name, extension, and a generic argument such as an asterisk)
3. Directory-tree backup (display a tree diagram and select files, or choose to back up the entire directory).

A version-based save means that ADSM stores a number of different versions of a saved file. The number of versions to be saved is defined by the ADSM administrator.

The backup function is commonly used for recovery from a disaster, user error, disk crash, or when performing a disk upgrade.

- Archive function

The archive function allows you to save a file that has a retention period specified in a number of days. Archive, too, can be specified by nominating a specific file (or using a generic argument such as an asterisk), or selecting from a directory tree (or everything in a subdirectory). Archive is used for:

- A more permanent storage of static objects
- Space management (because saved objects can be deleted from the PC)
- Software distribution (by centrally storing programs and files that can be retrieved by other users)
- Sharing data (by backing up on one client and restoring to others)

With ADSM/400 Version 2, you can specify that a file is to be deleted on the client after it has been archived on the AS/400 server.

- Restore function

The restore function recovers files that have been saved with a *backup* operation. There are three types of restore:

1. By file specification
2. By directory tree
3. By subdirectory path

The restore can occur to the client that performed the backup or to any other authorized client workstation.

- Retrieve function

The retrieve function recovers files that have been saved with an *archive* operation. Retrieval is based on one or more of the following factors:

- File specification (name or extension)
- The description associated with the archive operation
- Save date range
- Expiration date range

The retrieve occurs to the client that performed the archive or to any other authorized client workstation.

- Other Client Characteristics

Backup and Archive also offer the following features:

- Compression—To reduce network traffic and save server storage, data can be sent to the server in a compressed format.
- File filtering (include or omit from backup operations)—Selected files can be omitted from the backup so that client operating system files that cannot be restored are not saved. This filtering is managed by an include and exclude list on each client.
- Restore overrides—This feature restores files to subdirectories and file names different than the ones from which they were backed up.

9.1.3 Administrative Client Functions

The administrative client is in control of the whole ADSM environment to:

- Define workstations that are to be backed up
- Define the security setup for ADSM
- Administer the usage and setup of disk and tape
- Set up schedules for backups and enroll workstations with the ADSM server
- Define how long to retain files when they have been backed up to the AS/400 server and in how many versions
- Monitor schedules and events

The administrative client runs on a variety of PC operating systems such as OS/2, Windows 3.1 and above, Windows NT, AIX, Hewlett-Packard HP-UX, Sun Microsystems SunOS, or Solaris (all which provide a graphical user interface).

The administrative client comes as part of the ADSM server code for the AS/400. Using the AS/400 server administrative client, gives you a command-line, green screen interface (not a GUI interface).

9.1.4 Server Functions

The server has the same broad functions across all supported server platforms. It is the basic building block of ADSM and performs a number of tasks, which include:

- Storing client data in *volumes* within *storage pools* (disk and tape)
- Tracking files that have been backed up in the recovery log (a file on disk)
- Retaining rules about the backup and archive process (policies) in the ADSM database (another file on disk)
- Offering device support for client data storage, including to DASD, tape (8mm, QIC, 3490, and AS/400 reel devices), and optical (the AS/400 server platform does not support the use of optical for this function, but other server platforms do)

9.1.5 Total Disaster Recovery with ADSM/400 Version 2

On Version 2 ADSM/400 systems, ADSM has the ability to make a total disaster recovery plan in case your AS/400 server breaks down or AS/400 disks are damaged. There are several ways in which you can protect your AS/400 ADSM server. The simplest one is to let everything (storage pools, database, and recovery log, for example) stay in the AS/400 library QUSRADSM, which is the

default. Back up this library every day. This, however, requires that all user storage pools stay on AS/400 DASD, which could be an expensive solution. Many ADSM/400 installations use the option to reserve a certain amount of disk space for backed up user data. Then define a *tape storage pool* to which data is *migrated* when the disk area fills up. This setup requires a more complex backup strategy for the ADSM/400 server, such as:

- **Server database backup and recovery** is designed to help recover the ADSM database on the AS/400 server in the event of the loss of all or part of the database. Without the database, ADSM has lost track of which client files have been backed up and where they are stored. The server database can be backed up to tape with an ADSM/400 command. The tape should then be moved offsite.
- **Storage pool backup and recovery** is designed to help recover ADSM storage pools on the AS/400 server in the event of the loss of all or part of the ADSM storage pools. This function supports incremental backups and allows backups to take place while the server is operating and available to clients. Storage pools backed up can be both disk and tape storage pools.

As part of a disaster recovery plan, ADSM for AS/400 can identify removable media volumes containing storage-pool backup data to be sent to an offsite location. In turn, ADSM for AS/400 can identify which volumes that you may wish to return in the event of a loss of the physical ADSM server or the loss of a few volumes at the on-site location.

9.1.6 More ADSM/400 Information

ADSM/400 Version 2 runs under OS/400 V3R2, V3R7 and subsequent releases. Note that while ADSM for the AS/400 server is at Version 2, ADSM for other operating systems is available in Version 3.

Note: In a future release of OS/400, an OS/400 ADSM client will be made available. This means that with ADSM, user data from an AS/400 system can be backed up to another AS/400 system, a mainframe system, or a RISC/6000 system.

For more information about ADSM/400, see the following manuals (among others):

- *ADSTAR Distributed Storage Manager: General Information*, GH35-0131-02
- *ADSTAR Distributed Storage Manager for OS/400: Administrator's Guide*, SC35-0196
- *ADSTAR Distributed Storage Manager for OS/400: Administrator's Reference*, SC35-0197

For other information, such as the most current list of devices supported, see the ADSM Internet homepage:

<http://www.storage.ibm.com/storage/software/adsm/adsmhome.htm>

9.2 ADSM/400 and BRMS/400 Interoperability

BRMS/400 is IBM's premier offering for media management and performing save and restore of AS/400 objects. As both BRMS/400 and ADSM/400 input and output to and from tape, there are some situations where the two products overlap. For a detailed description of the interoperability of the two products, see the redbook *Setting Up and Implementing ADSTAR Distributed Storage Manager/400*, SG24-4460.

In ADSM/400 Version 2, ADSM/400 extends the tape interface so that ADSM offers a more flexible handling of tape operations. This can help you decide which tapes to restore from an offsite location in the event of a server failure.

Prior to ADSM Version 2, tape inventory functions are supported by the interface between ADSM/400 and Backup Recovery and Media Services/400 (BRMS/400). For tape inventory on these level systems, implement BRMS. ADSM/400 Version 2 works with tapes without the need for BRMS.

BRMS/400 and ADSM/400 are non-competing, stand-alone products. There are *no* interdependencies. BRMS/400 is complementary to ADSM/400 because it performs:

- Volume media management:
 - Initialization and enrollment
 - External tape volume labeling
 - Inventory tracking and reporting of tape media
 - Tape volume attribute and error tracking
 - Access security
 - Media expiration management
- Backup and recovery of ADSM/400 databases, ADSM/400 recovery logs, and ADSM/400 disk storage pools.

However, since ADSM/400 Version 2 offers functions for backup and recovery of the ADSM/400 database, recovery logs, and disk and tape storage pools, the need for BRMS/400 from an ADSM/400 point of view has been reduced. BRMS/400 is still needed for total management of your backup and tape environment for files outside of ADSM.

9.3 Backup, Recovery, and Media Services for AS/400 (BRMS)

Many customers with large systems or networked environments find it difficult to easily keep track of all removable media. BRMS for AS/400 facilitates centralized management of media by maintaining a consistent view of removable media, its contents, location, and availability across multiple AS/400 systems.

BRMS for AS/400 is a policy-oriented setup and execution of archive, backup, recovery, and other removable media related operations. Archive, backup, and recovery facilities enable the customer to establish how these operations are to be performed, from a high-level overview down to a file member level of control. BRMS provides the flexibility to allow custom programmed extensions, if necessary.

Media, whether used for backup or other operations, can be managed and traced in various ways (for example, by volume ID, type, content, location, container, and customer policies). Operation planning facilities assist the

customer in anticipating resources (devices, media, operational steps, and so on) needed to effectively and successfully complete the operation. Operations are guided, making it easier and less error-prone. Policy support throughout enables you to establish control and evaluation criteria on an operation, object or group, media, device, or system basis.

Note: BRMS code was “skip-shipped” at V4R2. This means that BRMS code does not need to be re-installed when other licensed program products are installed. The V4R1 version of BRMS runs on both OS/400 V4R1 and V4R2.

9.3.1 BRMS Recovery Report

One component of BRMS that can justify the cost of the entire product is the recovery report. After backups are complete, a full system recovery can be driven by the reports produced. The reports indicate:

1. What needs to be restored to the system
2. The order in which to restore the data
3. Which media contains the saves

Imagine what it takes to produce a report like that manually!

Menu driven recovery counts down the numbers of objects remaining to be restored, giving the user a clear indication of how the recovery is progressing. Aside from full-system recovery, individual objects, including database file members, can be restored easily using drill-down menu options, or by BRMS commands. The recovery report makes recovery planning a non-complicated, very manageable task.

9.3.2 BRMS User Exits and Message Handling

BRMS provides flexibility for controlling the AS/400 backup environment. Escape messages are sent if some listed items are not saved during a BRMS backup. BRMS messages signal an abnormal ending which starts up an alternate job on the job scheduler. These messages allow you to monitor BRMS save jobs more easily and correct potential backup exposures. By monitoring for these escape messages you can more quickly ascertain when a backup is not complete.

Customers who use job schedulers are aware that not all items are backed up using the job scheduler. To avoid an *ALTERNATE job to be kicked off automatically if the BRMS save does not save everything change the severity of the message itself.

This ensures the severity is under the limit for an abnormal end according to your job scheduler.

The CHGMSGD command can be used on the messages listed below. Read the text to determine if you want to retain any of these messages. For example, the BRM15A1 message for not saving vital BRMS recovery information is a message to retain at severity 70.

- BRM10A1—Library not completely saved (shipped at severity 70).
- BRM14A1—List not completely saved (shipped at severity 70).
- BRM15A1—BRMS Media Information vital for recovery not all saved (shipped at severity 70).
- BRM16A1—Backup completed with errors (shipped at severity 70). This message is sent at the end if any of the others are sent.

These messages are in the Q1AMSGF file in the QBRM library. If you have secondary languages, the messages may be in the QSYSnnnn language library where nnnn is the number of the secondary language.

If you are not certain how these messages affect your backups, let the default values remain. If you notice a conflict with alternate jobs being started, make changes as necessary.

9.3.3 BRMS Maintenance Activities—STRMNTBRM

In the interest of improving performance and recoverability, you can select BRMS maintenance activities. Users can split up parts of maintenance to run at various times, performing normal maintenance activities on a daily basis, while gaining the ability to perform single reporting options without needing to run all the cleanup, media auditing, and other tasks. This is done by selecting and deselecting the parameters on the STRMNTBRM command to:

- Execute BRMS cleanup
- Produce a volume session statistics report
- Print a version control report for media expiration
- Produce a recovery analysis report to view the recovery plan as if the only media available is at a specified location (such as offsite)
- Print a media information report
- Print a system audit media report indicating the changes and differences in the backup that could not be resolved and require user action

Note: The audit report notes the differences in the backup. It does not need to be produced every time maintenance is run unless a back-level QUSRBRM database has been restored to one system or a loss of QA1ANET records has occurred. Entries on this report indicate possible network system communication problems that should be investigated. These include:

- **RUNCLNUP parameter *YES or *NO:** Cleanup operations should be run regularly. However, if maintenance is run for a report-only reason, this parameter may be set to *NO.
- **RTVVOLSTAT parameter value *YES or *NO:** Volume session statistics are produced or not produced. If a site does not refer to these statistics, this option should be *NO whenever maintenance is run.
- **AUDSYSMED parameter values *NETGRP and *NONE:** Media information in a BRMS network is audited and differences resolved. A report is printed indicating the changes and differences that could not be resolved and which require user action. Ideally one system in the network runs the audit, making the run time for maintenance on the other systems shorter. This option does not need to be run every time maintenance is run, unless a back level QUSRBRM database has been restored to one system, or a loss of QA1ANET records has occurred. Having any entries on this report indicates possible network system communication problems that should be investigated.
- **PRTVSNRPT parameter values *YES, *NO or *EXPMED:** A version control report is or is not produced. *EXPMED indicates the report should be produced only if media expiration processing is performed. If your site never uses versioning, change the parameter to *NO. A site may only require this report on a weekly basis or less.
- **A multi-valued recovery locations parameter, RCYLOC:** This parameter depends on PRTRCYRPT(*RCYANL) or PRTRCYRPT(*ALL) to provide users

with the ability to qualify the recovery analysis report location. This allows users to view their recovery plan as if the only media available were at a specified location (such as offsite). A special value of *ALL is the default. The STRRCYBRM command also uses the *ALL parameter. Users might consider regularly running two or more recovery reports, based on different locations to cover various disaster possibilities (for example, loss of all on-site media).

9.3.4 BRMS BRMLOG

The BRMLOG report produced by running the Work with Move Policies command (WRKPCYBRM) tracks when the media monitor is enabled and disabled. Message BRM15A6 indicates that the media monitor has been disabled. Message BRM15A5 indicates that the media monitor has been enabled again.

The BRMLOG report time stamps agree with joblog time stamps for all V4R1 save and restore operations. Prior to V4R1, these values were written to the BRMLOG at the end of the save operation and showed that time stamp instead.

9.3.5 Media Classes—CPYMEDIBRM

Media classes with the same name but different densities when using OPTION(*FROMFILE) are diagnosed with the CPYMEDIBRM command. The differences are listed in the QP1AEN report. The user should change media class attributes as appropriate. On earlier BRMS versions, these differences in incoming media classes were ignored, resulting in selection of inappropriate media by BRMS.

9.3.6 Network Time Zone Synchronization

BRMS allows systems in a network to be in different time zones and yet maintain accurate time adjusted records of all media information and saves. This time zone support is available through PTFs in V3R2/6/7 but included in BRMS V4R1 and later code. All systems in the network *must* contain the new code in order to implement various time zones. The network must not contain a system at V3R1 if various time zones are represented. If various time zones are not implemented, no changes are necessary.

If various time zones already exist on an AS/400 network and BRMS is being newly implemented, no special tasks are necessary except to ensure that the necessary supporting PTFs are applied to pre-V4R1 systems before you follow the instructions for creating your BRMS network. This is outlined in the *Backup Recovery and Media Services for AS/400*, manual SC41-4345.

We recommend that you run the INZBRM *NETTIME command on BRMS systems prior to V4R1 to keep all system clocks in sync on a BRMS network. This command resets the QTIME system value on each system in the network so that it matches the QTIME of the system running the command.

Some sites run INZBRM *NETTIME automatically from a job scheduler. When utilizing job schedule entries, ensure that such a job-schedule entry is removed before making any changes to system times.

Recommendation

If differing time zones are implemented, do not run the INZBRM *NETTIME command on your networked systems until time zone differences are accounted for.

Considerations when system times are changed on pre-existing BRMS networks include:

If you are setting clocks forward in time:

1. Ensure that all BRMS activity (such as saves, moves or maintenance) is ended on all systems.
2. Ensure that the BRMS subsystem Q1ABRMNET is ended on all systems in the network.
3. Set the time forward on any systems requiring a gain in time.
4. Restart the subsystem after all clocks have been set.

If the time is to be set backwards on any system in the network:

1. Ensure that all BRMS activity, such as saves, moves, or maintenance, is ended on all systems.
2. Ensure that there are no updates waiting to synchronize across the network.
To do this:
 - Run DSPPFM QUSRBRM/QA1ANET on each system in the network. If any members exist, updates are waiting to synchronize.
 - Wait a few minutes to allow the synchronization job to send the records again. Retry the DSPPFM command.
 - If there are no changes to the file after several minutes, there could be a communications problem. Correct any communications problems, and retry the DSPPFM command. Consult the AS/400 Support Line if necessary.
 - When there are no longer any members in the QA1ANET file, all updates have been sent.
3. End the BRMS subsystem Q1ABRMNET on each system in the network.
4. If you are setting a system back “n” hours, power the system down, wait “n” hours and manually IPL the system and reset the system time. This helps you avoid potential problems with duplicating time stamped records in BRMS and other system journals. Duplicate time stamps can cause unpredictable results for BRMS media management.
 - If it is not possible to power the system down for “n” hours to make the time change, ensure that no BRMS activity occurs for “n” hours on the *entire* network.
 - If several systems are being set backwards to various times, such as 5PM, 6PM, and 7PM, use the system that is set back the furthest in time to represent the value for “n” in the preceding description.
 - Be aware that other system journals and logs, such as the QHST logs, may contain confusing or inaccurate information as a result of keeping the system up when setting the clock backwards in time.

Note: To ensure the current QHST log version is up to date, use DSPLOG with a fictitious message identifier, such as:

DSPLOG MSGID(###0000)

DSPLOG of a fictitious message ID does not display any messages, but the log version physical file is made current.

9.3.7 Large Tape File Sequence Numbers for Non-Save and Restore Operations

Due to the increased storage capacity of many of the tape drives in current technology, BRMS supports up to 65 535 sequence numbers. In releases prior to V4R1, the maximum sequence number is 9 999.

Note: The support for up to 65 535 sequence numbers per tape is for non-save and restore operations only. This affects BRMS users using the new append options on DUPMEDBRM. DUPMEDBRM does not support larger sequence numbers for BRMS save operations, such as those indicating APPEND *YES, because of OS/400 save-and-restore restrictions.

The QUSRBRM library files affected are QA1AHS and QA1AOD. QA1AHS is sometimes shared between systems. Therefore, previous releases experience a restriction when it comes to storing media information on a shared inventory system. On systems prior to V4R1, the field containing the sequence number information is not large enough to hold values over 9 999. Sequence numbers do not appear properly when such a system is receiving *LIB level information from the V4R1 systems.

Note: Under WRKPCYBRM *SYS, the Change Network Group option listed in the System Policy indicates if a system is sharing media information with other systems in the network by using the *LIB special value.

9.3.8 Shorten the Time to Save with USEOPTBLK

The Use Optimum Block size attribute is added to SAVxxxBRM and SETMEDBRM commands in V4R1 to allow CL program users with supported devices to use the larger block sizes, for example the 3590 and 3570 tape drives. Special values of *YES, *NO, and *DEV are supported. *DEV indicates that the BRMS device attribute (found in the listing of WRKDEVBRM) should be referenced.

Backup policy, archive policy, and control group attributes controlling the USEOPTBLK option are also provided. When values other than *DEV (the default) are specified on policies or control group attributes, these take precedence over current settings under WRKDEVBRM listings.

Users of the 3590 Magstar MP tape library experience the most performance increase by specifying *YES for USEOPTBLK. If USEOPTBLK is set to *YES for the 3570 Magstar tape library, there may be a slight performance improvement.

Refer to Section 5.6, "Use Optimum Blocking for Save and Restore" on page 57, for further information on USEOPTBLK.

9.3.9 SAVSYSBRM Command Support for OMIT Parameter

Similar to the changes made to the SAVxxx and RSTxxx commands, objects can be omitted from a BRMS save by specifying the OMIT parameter. Objects to be omitted include:

- *CFG
- *SECDTA

- *NONE

This allows configuration objects and security objects to be omitted from the SAVSYS, which is performed using the SAVSYSBRM command.

Recommendation

If you OMIT *CFG or *SECDATA for a shorter save time or less media use, ensure the *CFG and *SECDATA objects are periodically saved separately.

9.4 WRKASP Utility

The WRKASP utility for OS/400 PRPQ is available for V4R2, V4R1, V3R7, V3R6, and V3R2. This PRPQ provides planning and management functions to use and manage multiple ASPs. Product and PRPQ numbers for the supported releases are as follows:

- 5799-GDH PRPQ P84308 for V4R1 and V4R2
- 5799-FZG PRPQ P84300 for V3R6 and V3R7
- 5799-FZF PRPQ P84301 for V3R2

An auxiliary storage pool (ASP) is a logical grouping of disk units. ASPs provide a means of isolating objects on a specific disk units. If the system experiences a disk unit failure that requires its replacement, recovery is required only for the library or objects in the ASP containing the replaced disk unit.

When a user ASP is created (using DST), a physical set of disk units is logically grouping together. User ASPs contain either entire libraries or associated objects only. Associated objects can be journals, journal receivers, and save files.

Isolating libraries and associated objects in a user ASP protects them from disk failures in other ASPs (either the system or other use ASPs). For example, if a disk failure causes data to be lost in user ASP03, data in the system ASP and other user ASPs remain intact and do not have to be restored.

WRKASP allows you to work with ASP libraries, overflowed objects, and save files, journals, and journal receivers using a menu driven interface. Included in the WRKASP utility are seven CL commands, which include:

WRKASP	Starts the main WRKASP menu.
CHGASPDSC	Allows unique ASP descriptions within WRKASP. This is useful when objects are categorized by application or the type of object (for example, all spool files, or all billing application data).
CPYLIBASP	Duplicates a library from one ASP to another.
MOVLIBASP	Moves a library from one ASP to another.
PRDBDP	Prints database dependencies for a library. This command provides similar information as that produced by the WRKJRNA command or ANZDBF command in the Performance Tool/400 product.
PRTASPLIB	Prints a report of ASP contents and save dates.
SAVLIBASP	Saves the contents of a library based on user ASPs.

WRKASP simplifies the ASP planning and implementation process, as well as the day-to-day management and maintenance of user ASPs. One of the most important features of the tool is the ability to move libraries across ASPs and reset overflowed objects without interrupting normal operations by preserving private authorities to system objects, as discussed in the sections that follow.

9.4.1 ASP Monitoring

Monitoring of an ASP status is done from a single panel or menu. The user quickly determines how much disk space is used, the status of the disk protection for that ASP, and whether it is in an overflow status or not. The following display shows an example disk configuration for user ASP02.

Display Disk Configuration by ASP							
ASP number : 02							
Unit	Serial Number	Type	Model	Size	% Used	Protection Type	Overflow Status
016	00-OCA9322	6607	0072	3670	32.1	DPY	Active
017	00-OCA6391	6607	0072	3670	32.1	DPY	Active
018	00-OCC0727	6607	0072	3670	32.1	DPY	Exposed
019	00-ODF4248	6607	0072	3670	32.1	DPY	Active
020	00-OCC0346	6607	0072	3670	32.1	DPY	Active
021	00-0BF6449	6607	0072	3670	33.5	MRR	Exposed
022	00-OAC1776	6607	0072	3670	33.5	MRR	Active

Figure 30. Disk Configuration by ASP Showing ASP02

9.4.2 ASP Configuration

There are three types of ASPs on the AS/400 system. They are:

- SYS** For the system ASP (always ASP01)
- LIB** For library-based ASPs
- OBJ** For object-based ASPs

The system ASP contains IBM supplied functions, such as the Licensed Internal Code, the operating system (including library QSYS), document library QDOC, libraries QUSRSYS and QGPL and configuration objects. Some user objects, such as programs, files, and libraries, can be contained in the system ASP.

If no user ASPs exist, all configured DASD units are allocated to the system ASP, and all libraries and objects are stored in the system ASP.

Library-based user ASPs contain entire libraries, including the objects that reside in that library. Objects are designated for a specific user ASP by specifying a value of 2 through 16 in the ASP parameter of the Create Library (CRTLIB) command.

Object-based user ASPs do not contain libraries. Object-based user ASPs are used to contain journals, journal receivers and save files, whose associated library is stored in the system ASP. Again, objects come to reside in a specific

user ASP when a value of 2 through 16 is specified on the create command, for example:

```
CRITSAVF FILE(file-name) ASP(02) TEXT(Save File in User ASP 02)
```

The protection type can be *DPY* for device parity protection, *MRR* for DASD mirroring, *CKS* for mixed ASP protection (mirroring with device parity), or blank for no protection.

The protection status can be *active* when protection for all mirrored and device parity DASD in this ASP is active. The *exposed* status indicates that mirrored protection for a unit in the ASP is currently suspended or a device parity drive has a problem. The *resume* status indicates that a mirrored unit has been suspended but is now resuming, or a device parity protected unit has been repaired and the unit is being rebuilt.

9.4.3 ASP Management

WRKASP provides an effective view and method to manage ASPs, such as:

- Locate and reset overflowed objects

You have the ability to reset overflowed objects without having to manually save, delete, restore, or IPL the system.

- Locate database networks, invalid objects, and journal dependencies

This report can be used to determine if cross library database networks exist (that is, if there are logical files in a different library than the physical) and if the ASP has enough capacity for the move.

The following display is the main panel of the WRKASP utility:

```

                                Work with Auxiliary Storage Pools
                                System:  SYSTEMXX

Type options, press enter.
  5=Display disk status      6=Print libraries      9=Save ASP libraries
 11=Save change contents   12=Work with contents  13=Change description
 22=Print ASP Analysis

Opt  ASP  ASP description          Type  --Protection--  Over
      1  System ASP 1 - OS/400  SYS   Type  Status    flow
      2  User ASP data and libraries  LIB  DPY           YES
      3  Test data              LIB  MRR  Active    NO
      4  Programs in development  LIB  DPY           NO
      5  History data           LIB           NO
      6  Source files           LIB  CKS           YES
      7  Save files             OBJ           NO
      8  Spool files            LIB  MRR  Active    NO
      9  Payroll data and libraries LIB  MRR  Active    NO

Command
===>

```

Figure 31. Work with Auxiliary Storage Pools

From this display, option 12 *Work with contents* takes you to the following display:

```

Work with contents (Libraries)

ASP number . . . . : 02           Position to . . . .
Available ASP space: 42073014

Type options, press enter.
3=Copy to ASP   5=Display  6=Print DB   8=Display description
11=Move to ASP 12=Work with library 13=Analyze library

Opt  Library      Description                Size  Obj  DB  Jrn
-----
DMTLIB  Dental Application Library 26836 No   No  No
SCRATCH Testing scratch pad       77   No   No  No
QUERIES COLLECTION - created by SQL 888  No   Yes Yes
AAALIB  Human Resources initiation 143  No   No  No
DSPUS   Display user space library 167  No   No  No
AAALIB  Human Resources application 420986 No   No  Yes
SLCSXP  COLLECTION - created by SQL 200560 No   Yes Yes

More...

Command
===>
F3=Exit  F5=Refresh  F9=Retrieve  F12=Cancel  F13=InfoSeeker
F15=Work with save file and journal objects

```

Figure 32. *Work with Contents (Libraries) Display*

By looking at the **Dependencies** column, you can immediately see if you should analyze one or more libraries further for database dependencies. In our example, the QUERIES, AAALIB, and SLCSXP libraries have file dependencies.

Option 6, **Print DB** gives a printed analysis of the database dependencies for each library.

- ASP-to-ASP library move and copy

This option provides the ability to move or copy a library from one ASP to another with a single command and without losing private authorities. When system commands are used to save a library, delete it and restore it to another ASP. Private authorities are lost and need to be restored (using RSTAUT) or recreated (using EDTOBJAUT).

For more information about the WRKASP utility, refer to *WRKASP Utility User's Guide PRPQ*, SC41-0652.

9.5 Print System Information Tool

The PRTSYSINF tool was introduced as part of the Upgrade Assistant when CISC to RISC migrations were initiated. Beginning with V4R2, the PRTSYSINF command is part of OS/400.

A record of the contents of your system, such as how your system is customized and what libraries it contains, is important to your upgrade success. The Print System Information (PRTSYSINF) command prints system information that should be maintained for disaster recovery and system verification purposes. The information helps you:

- Plan your upgrade procedures.
- Evaluate the success of moving information.
- Perform disaster recovery, if necessary.

The PRTSYSINF command can also be useful in your daily operations even if an upgrade is not planned.

The PRTSYSINF command generates the following reports (the corresponding OS/400 CL command is in parentheses):

- A Library Backup list with information about each library in the system, which backup schedules it is part of, and when it was last backed up (DSPBCKUPL *LIB)
- A Folder Backup list with the same information for all folders in the system (DSPBCKUPL *FLR)
- A list of all the system values (DPSYSVAL)
- A list of network attributes (DSPNETA)
- A list of edit descriptions (DSPEDTD)
- Display PTF Details list (DSPPTF)
- A list of reply entries (WRKRPYLE)
- A report of access path relationships (DSPRCYAP)
- A list of service attributes (DPSVRA)
- A list of network server storage space (DSPNWSSTG)
- A report showing the power on/off schedule (DSPPWRSCD)
- A list for all hardware features on your system (DSPHDWRSC)
- A list of distribution queues (DSPDSTSRV)
- A list of all the subsystems defined on your machine (DPSBSD)
- A list of the IBM software licenses installed on your machine (DPSFWRSC)
- A basic list of journal object descriptions for all journals on the system (DSPOBJD)
- A report showing journal attributes for all journals on the system (WRKJRNA)
- A report showing cleanup operation (CHGCLNUP)
- A basic list of all user profiles on the system (DSPUSRPRF)
- A report of all job descriptions on the system QDFTJOB (DSPJOB)

All the commands listed above are part of OS/400 and can be run individually, but the PRTSYSINF command ties them together neatly.

Note: As some of the above commands require that you have *SECOFR authority, you also need *SECOFR authority to run the PRTSYSINF command.

RTVSYNINF is used to retrieve the same values into save files, user spaces, and data areas of a library of your choosing for further analysis.

UPDSYSINF can selectively update the files that RTVSYNINF created. with a current status of the edit descriptions, network attributes, reply list entries, service attributes, service providers, and system values on the system. These tools are very useful to help the system administrator keep track of how the system is configured and customized.

9.6 OS/400 Job Scheduler (Part of the Operating System)

The OS/400 job scheduling function inherent in the operating system allows for time-dependent scheduling of AS/400 batch jobs. You can schedule jobs to be released from the job queue at a particular time, or you can use a job schedule entry to submit your job to a job queue automatically at the time you indicate.

Job scheduling allows you to control the date and time that a batch job is submitted or becomes eligible to start from a job queue. This flexibility helps you balance the workload on your system.

9.6.1 Work with Job Schedule Entries

The Work with Job Schedule Entries display shows scheduling information for job schedule entries, the job queues to which the jobs will be submitted, the date and time of the last successful submission, the user profile that added the entry to the job schedule, and a brief description. You can work with the job schedule entries shown in the list by entering an option number next to the entry and pressing Enter.

```
Work with Job Schedule Entries                                SYSTEMXX
                                                           01/22/98 20:27:22

Type options, press Enter.
 2=Change  3=Hold  4=Remove  5=Display details  6=Release
 8=Work with last submission 10=Submit immediately

-----Schedule-----
Opt  Job          Status Date      Time      Frequency  Recovery  Next
-----
  _  ADMRESSAV     HLD  *ALL     23:30:00  *MONTHLY  *SBMRLS   02/07/98
  _  QSECIDL1     SCD  *ALL     01:00:00  *WEEKLY   *SBMRLS   01/26/98
  _  RTVDSKINF    SCD  *SAT     00:30:00  *WEEKLY   *SBMRLS   02/27/98
  _  TTLMOM       SCD  *ALL     00:05:00  *WEEKLY   *SBMRLS   01/22/98
  _  SLTSXP       SCD  *ALL     00:05:00  *ONCE     *SBMRLS   02/14/98

                                           Bottom

Parameters or command
===>
F3=Exit  F4=Prompt      F5=Refresh  F6=Add  F9=Retrieve
F11=Display job queue data  F12=Cancel  F17=Top  F18=Bottom
```

Figure 33. Work with Job Schedule Entries Command

To add a job schedule entry, select F6 to add.

Add Job Schedule Entry (ADDJOBSCDE)

Type choices, press Enter.

Job name _____ Name, *JOB

Command to run _____

Frequency _____ *ONCE, *WEEKLY, *MONTHLY
Schedule date, or *CURRENT Date, *CURRENT, *MONTHSTR...
Schedule day *NONE *NONE, *ALL, *MON, *TUE...
+ for more values

Schedule time *CURRENT Time, *CURRENT
Relative day of month 1 *LAST, 1, 2, 3, 4, 5
+ for more values

Save *NO *NO, *YES

More...

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 34. Add Job Schedule Entry Command

Job scheduling information is contained in the job schedule system object named QDFTJOBSCD. This object cannot be deleted, renamed, copied to another library, or duplicated. It must always reside in the QUSRSYS library.

Should the QDFTJOBSCD object ever become damaged, the system replaces it with an empty object, which means that none of your scheduled jobs will run. You can restore the QDFTJOBSCD object from your backup during production. You do not have to be in a restricted state. You can also restore the job schedule object from another system and use Change Job Schedule Entry (CHGJOBSCDE) command to modify it for your system.

An empty job schedule object can be the result of the QDFTJOBSCD user-modified object being over-written, such as when library QUSRSYS is replaced in a release update operation. Therefore, it is important that your backup plan include this object. Keep in mind that this object must *always* reside in the QUSRSYS library.

Note: Another reason for scheduled jobs not executing is that the QJOBSCD job is not running. The system QJOBSCD job starts automatically during IPL. It remains active until the system is powered down. If it inadvertently ends, a SAV or RST operation of the QDFTJOBSCD object in QUSRSYS reactivates normal job schedule functions. However, if the SAV or RST is run in restricted state, job scheduling does not activate until normal operations are resumed.

Additional considerations apply when restoring the job-schedule object, if it is restored from a system that uses a different calendar format (for example YYMMDD date format on a MMDDYY system). Failure to reconcile the difference in a date format can cause unpredictable job scheduling.

Further information about using the job scheduler function can be found in the *Work Management Guide*, SC41-5306.

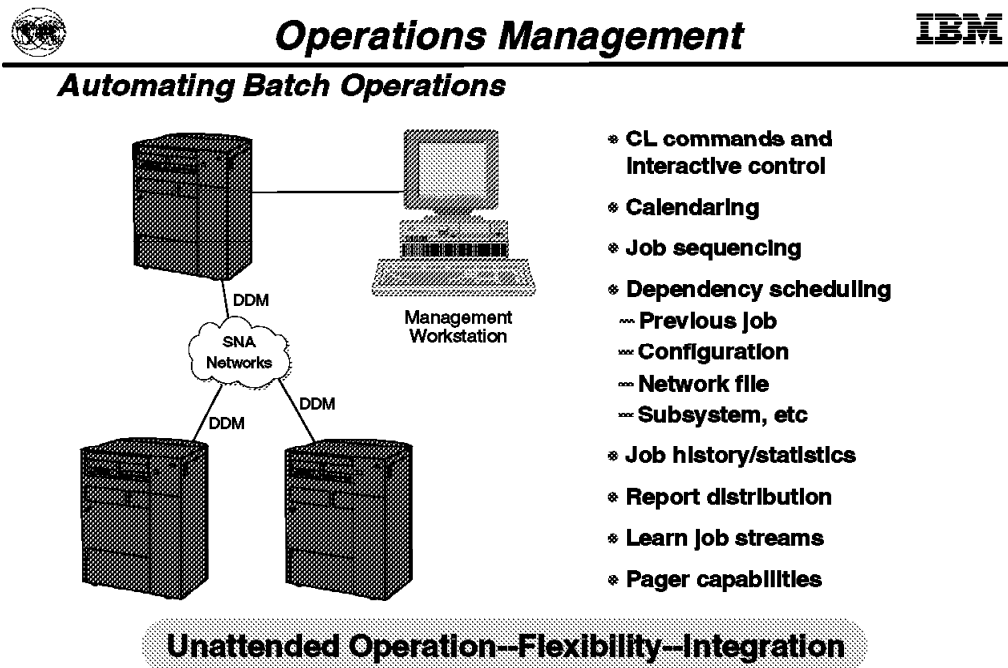
9.7 IBM Job Scheduler for OS/400

It is not enough that the system is up and running when you talk about system availability. From an end user's point of view, it is equally important that their applications run as they are expected to, so that the daily work can be accomplished. The functions described here help you ensure that applications run at the time that they are meant to run and that they complete normally.

9.7.1 Job Scheduling and Availability

If you want more advanced functions than what the job scheduler in OS/400 provides—such as job dependencies, where batch job two is only run if batch job one completes successfully—look at IBM Job Scheduler for OS/400. This product takes automation further than what is provided by the job scheduling function provided in OS/400.

IBM Job Scheduler for OS/400 facilitates unattended operations and provides a highly comprehensive, full-function job scheduling and report distribution system on the AS/400 system. It handles conditions such as previous job completion, completion code checking, date file availability, and operator input.



AACSVLPRE-19

IBM Corporation 1996

08/96

Figure 35. IBM Job Scheduler/400 Features

IBM Job Scheduler for OS/400 can be set up with data currently used by the date and time scheduler found in OS/400. When OS/400 scheduling entries are converted to IBM Job Scheduler for OS/400, you enhance your scheduling capabilities through the use of calendaring dependency scheduling, plus other features provided by IBM Job Scheduler for OS/400.

The uses of IBM Job Scheduler for OS/400 are numerous (refer to Figure 35). Job Scheduler for OS/400 provides a key function that moves AS/400 system operations even closer to a total “lights-out,” automated set of operations.

Highlights include:

- Automation
- Batch job-stream management
- Forward planning and production forecasting
- Full calendaring of operations
- Dependency scheduling

Any batch-capable function can be scheduled. The product interfaces with other IBM systems management tools for the AS/400, giving you functions such as:

- Enhanced scheduling of PTF orders through IBM SystemView System Manager/400
- Scheduling and maintaining backups by combining Job Scheduler for OS/400 with BRMS/400

The functions of Job Scheduler for OS/400 are available to run on a single system or across multiple systems in a network.

Additional functions include:

- Job reporting
- Repetitive distribution
- Capture existing job streams as they are submitted
- Pass job parameters using the local data area
- Operator notification using paging support

Note: Job Scheduler for OS/400 was developed and is maintained by MBA, Inc. Refer to the following URL for more information on MBA:

<http://www.mbainc.com/new/>


9.8 IBM SystemView System Manager for AS/400

The SystemView System Manager for AS/400 (SM/400) licensed program is part of the integrated offering Operations Control Center/400, which includes Managed System Services/400 (MSS/400). SM/400 is enhanced to integrate with SNMP management products, such as NetView for AIX. An SNMP manager can monitor for alerts, obtain system information and execute remote commands if the AS/400 system is to be managed from an SNMP platform. The change management functions support the integrated file system.



Software Distribution

- * **Distribute to**
 - Change Control Servers
 - Other focal points
 - Other AS/400 systems
- * **Plan distributions**
- * **Schedule activities**
- * **Immediate activities**
- * **Track distribution**
- * **Install software**
- * **Maintain software**



AACSVLPRE-15

IBM Corporation 1996

08/96

Figure 36. SystemView System Manager and Managed Systems Services

SM/400 provides central site control for:

- Remote AS/400 problem management
This includes remote problems analysis, comparison with existing available PTFs, automatic distribution of found PTFs, and a single connection to IBM or an independent software vendor for an electronic reporting of new problems.
- Central site packaging of independent software vendor (ISV) applications for AS/400 Licensed Program management support
This enables ISV applications to receive the same system support as IBM licensed programs.
- Central site distribution and change management support for remote AS/400 systems using MSS/400, remote RISC/6000 systems using NetView DM/2, and remote Novell NetWare Servers using NVDM for NetWare

SM/400 permits the central site AS/400 system to define, schedule, and track software distribution (change management) requests sent to systems with MSS/400, NetView DM/2, NetView DM/6000, or Novell NetWare installed. These change management requests include sending, receiving, and deleting files, programs, other AS/400 objects (libraries, save files, message files, documents, folders, PTFs, and so on) and non-AS/400 (OS/2 and RISC/6000) files, programs, and software.

Running programs, installing software, applying PTFs, and re-performing IPL can be scheduled to run automatically on the remote system. The remote system can forward the results of all change requests to the central site SM/400 system for tracking.

When only AS/400 systems are connected, the Operations Control Center/400 provides a significant set of automated operations. However, OCC/400 does not provide the *real-time* monitoring and automated action for the entire AS/400 operating environment.

- Sending AS/400 commands to remote AS/400 systems using MSS/400 without signing on

This support is intended for unplanned operations to be performed on one or more remote AS/400 systems, such as deleting a particular file or library that is no longer in use.

SM/400 includes a *graphical interface* for a network operator to graphically monitor and manage a network of systems. The change management functions are enhanced to provide support for the Integrated PC Server.

Note: SM/400 allows IBM licensed programs to have options that are at a different version, release, and modification level than the base product. The Copy PTF Save command (CPYPTFSAVF) removes its assumption for each product that the version, release, and modification of the option matches the base.

9.9 IBM SystemView Managed System Services for AS/400

The SystemView Managed System Services for AS/400 (MSS/400) licensed program is part of Operations Control Center/400, which includes SystemView System Manager for OS/400, as discussed in Section 9.8, "IBM SystemView System Manager for AS/400" on page 128. MSS/400 enables an AS/400 system to be managed from a central site running either:

- S/390 NetView Distribution Manager for MVS (Release 5 or later) for MVS-based networks
- SystemView System Manager for OS/400 (Version 3 Release 1 or later) for AS/400-based networks

The central site defines, schedules, and tracks software distribution (change management) requests sent to the AS/400 system with MSS/400 installed. Change management requests include sending, receiving, and deleting AS/400 files, programs, and other objects (libraries, save files, message files, documents, folders, PTFs, and so on).

AS/400 objects can be sent directly to or received from AS/400 libraries or through the local AS/400 distribution repository.

Running programs, installing products, applying PTFs, and re-performing IPL can be scheduled to run automatically under MSS/400 control. MSS/400 forwards the results of all change requests to the central site for tracking.

The capability for the central site to define, schedule, and run these change requests one time or repetitively significantly enhances unattended operation of remote AS/400 systems.

To make it easier to distribute software and manage changes for clients attached to AS/400 systems, MSS/400 uses a change control server function (SBCS only). This enables unaltered software distribution and installation for clients with NetView Remote Operations Agent/400 software (5733-165) installed. Support is

available for OS/2, Windows 3.1, DOS, UNIX, AIX, Windows NT, and Apple Macintosh** Clients.

MSS/400 also supports the unscheduled running of AS/400 commands issued by the central site, without having to first sign on to the AS/400 system with MSS/400. Printed output from these commands can be returned to the central site that issued the command.

9.9.1 Scenario of Using SM/400 with MSS/400

A software development department often creates a new release of the software package used for daily business operations. Each release has to be packaged and distributed to all other AS/400 systems in the network, where it will be installed from the central site system. To do this:

1. Create a product package.
2. Distribute the package to the remote systems.
3. Install the package on the remote systems.

SystemView System Manager/400 can treat user defined products as AS/400 licensed programs. You can define modification levels and PTFs for your own program packages.

9.10 Automating Message Management

As a system grows in size and numbers even up to clustered AS/400 systems, managing operational messages becomes an additional task for system administrators. The need to consolidate and automate messages becomes a priority so that unnoticed potential problems do not become “real” problems. Message management eases the burden on the operations staff, allowing them to focus on other issues. Tips for managing messages described in this section include:

- Using the QSYSMSG message queue
- Break handling program for message monitoring
- Setting a system reply list for an automatic reply of messages
- Operational assistant for operational management

9.10.1 QSYSMSG Message Queue

As AS/400 systems grow, so does the number of messages in the QSYSOPR message queue. They grow so much that many operators feel overwhelmed and miss important messages. OS/400 provides the ability to mirror critical system messages into an optional message queue named QSYSMSG in library QSYS. Once QSYSMSG is created, critical messages (such as DASD problems and security violations) are sent to QSYSMSG where the errors are more visible.

9.10.1.1 Creating the QSYSMSG Message Queue

The QSYSMSG queue does not get created automatically. Enter the following command to create it:

```
CRTMSGQ QSYS/QSYSMSG TEXT('Optional MSGQ to receive critical messages')
```

Using the QSYSMSG queue allows a user-written program to gain control when selected messages are sent. The program you write should be a break-handling program. An example of which is provided in Section 9.10.2, “Break Handling Program (User-Exit Program)” on page 132.

Note: Do not create the QSYSMSG queue unless you want it to receive specific messages. Once the QSYSMSG message queue is created, all the specific messages are sent to QSYSMSG without any further action from the operation.

9.10.2 Break Handling Program (User-Exit Program)

In addition to dealing with message volumes, all message-queue monitoring may be automated. There are products on the market that offer this, or you may create your own message-queue monitoring program, known as a Break Handling Program.

9.10.2.1 Creating a Break Handling Program

A break-handling program is one that is called automatically when a message arrives at a message queue in *BREAK mode. Both the name of the program and the break delivery mode must be specified on the Change Message Queue (CHGMSGQ) command.

The program must run a Receive Message (RCVMSG) command to receive messages. Parameters are passed to the user-defined program. The parameters identify the message queue and the message reference key (MRK) of the message causing the break. If the break-handling program is called, it interrupts the job in which the message occurred and runs. When the break-handling program ends, the original program resumes processing.

The following program is one example of a break-handling program.

```
PGM PARM(&MSGQ &MSGLIB &MRK)
DCL VAR(&MSGQ) TYPE(*CHAR) LEN(10)
DCL VAR(&MSGLIB) TYPE(*CHAR) LEN(10)
DCL VAR(&MRK) TYPE(*CHAR) LEN(4)
DCL VAR(&MSG) TYPE(*CHAR) LEN(75)
RCVMSG MSGQ(&MSGLIB/&MSGQ) MSGKEY(&MRK) MSG(&MSG)
.
.
.
ENDPGM
```

After the break-handling program is created, run the following command to link it to the QSYSMSG message queue.

```
CHGMSGQ MSGQ(QSYS/QSYSMSG) DLVRY(*BREAK) PGM(program-name)
```

Notes:

1. When a message queue is in break mode, any message on the queue calls the break-handling program. When messages are handled, remove them from the message queue.
2. The procedure or program receiving the message should be coded with a wait time of zero. This is done with the Receive Message (RCVMSG) command.

In the following example of a break-handling program, messages that are normally written to QSYSOPR, are handled from QSYSMSG instead.

```

BRKPGM:   PGM (&MSGQ &MSGQLIB &MSGMRK)
          DCL &MSGQ TYPE(*CHAR) LEN(10)
          DCL &MSGQLIB TYPE(*CHAR) LEN(10)
          DCL &MSGMRK TYPE(*CHAR) LEN(4)
          DCL &MSGID TYPE(*CHAR) LEN(7)
          RCVMSG MSGQ(&MSGQLIB/&MSGQ) MSGKEY(&MSGMRK) +
            MSGID(&MSGID) RMV(*NO)
          /* Ignore message CPA5243 */
          IF (&MSGID *EQ 'CPA5243') GOTO
ENDBRKPGM
/* Reply to forms alignment message */
IF (&MSGID *EQ 'CPA5316') +
  DO
    SNDRPY MSGKEY(&MSGMRK) MSGQ(&MSGQLIB/&MSGQ) RPY(I)
  ENDDO
/* Other messages require user intervention */
ELSE CMD(DSPMSG MSGQ(&MSGQLIB/&MSGQ))
ENDBRKPGM: ENDPGM

```

Note: This program cannot open a display file if the interrupted program is waiting for input data from the display.

In the example, if a CPA5316 message arrives at the message queue while the DSPMSG command is running, the DSPMSG display shows the original message causing the break and the CPA5316 message. The display waits for the operator to reply to the CPA5316 message before proceeding.

Note: When this program is used, the display station user does not need to respond to the messages:

CPA5243 (Press Ready, Start, or Start-Stop on device &1)

and

CPA5316 (Verify alignment on device &3)

The procedure within a user break-handling program may need suspend and restore procedures while the message handling function is performed. The suspend and restore procedure is necessary only if:

- A procedure in the break-program shows other menus or displays.
- The break-program calls other programs that may show other menus or displays.

The following example outlines the user procedure and display file needed to suspend and restore the display. RSTDSP(*YES) must be specified when the display file is created.

```

A   R SAVFMT           OVERLAY  KEEP
A*
A   R DUMMY           OVERLAY
A
A           KEEP
A           ASSUME
A   DUMMYR   1A  12DSPATR(ND)

```

```

PGM PARM(&MSGQ &MSGLIB &MRK)
DCL VAR(&MSGQ) TYPE(*CHAR) LEN(10)

```

```

DCL VAR(&MSGLIB) TYPE(*CHAR) LEN(10)
DCL VAR(&MRK) TYPE(*DEC) LEN(4)
DCLF FILE(UDDS/BRKPGMFM)
SNDF RCDfmt(SAVFMT)
CALL PGM(User's Break Program)
SNDF RCDfmt(SAVFMT)
ENDPGM

```

If you do not want the user-specified-break-handling program to interrupt the interactive job, submit the program to batch. You can do this by specifying a break-handling program to receive the message and perform a SBMJOB. SBMJOB calls the current break-handling program and passes any parameters necessary. Control is then returned to the interactive job which continues normally.

Refer to the *AS/400 CL Programming Guide*, SC41-5721-01, and "Break Handling exit program" in the *System API Reference*, SC41-5801-01.

9.10.3 System Reply List

Since all messages do not need to be reviewed, operator review and response to repetitive messages can be eliminated. Use the system message reply list to automatically reply with a specified default to inquiry messages.

Each entry in the system reply list specifies a message identifier, and the reply taken when that message is sent as an inquiry message.

The Work with System Reply List Entries display shows you a list of message identifiers and the replies to be sent when the system reply list is used.

Note: To cause the system reply list to be used, set the inquiry message reply parameter INQMSGRPY to *SYSRPYL for the job.

To see what reply list entries are on your system, use the PRTSYSINF command, as discussed in Section 9.5, "Print System Information Tool" on page 123.

9.10.4 Operational Assistant

Operational assistant is, among other things, an automated method of clearing old messages from message queues. The AS/400 Operational Assistant (ASSIST) menu simplifies some of the common user tasks, such as working with printer output, jobs, messages, and changing your password.

In addition, users with proper authority are given options to manage or customize the system, such as:

- Checking the status of the system
- Cleaning up objects
- Powering on and off
- Enrolling users in the directory
- Changing some system options
- Collecting disk space information


```
ASSIST                AS/400 Operational Assistant (TM) Menu                System:  SYSTEMXX
To select one of the following, type its number below and press Enter:

    1. Work with printer output
    2. Work with jobs
    3. Work with messages
    4. Send messages
    5. Change your password

    10. Manage your system, users, and devices
    11. Customize your system, users, and devices

    75. Information and problem handling

    80. Temporary sign-off

Type a menu option below
—
F1=Help  F3=Exit  F9=Command line  F12=Cancel
```

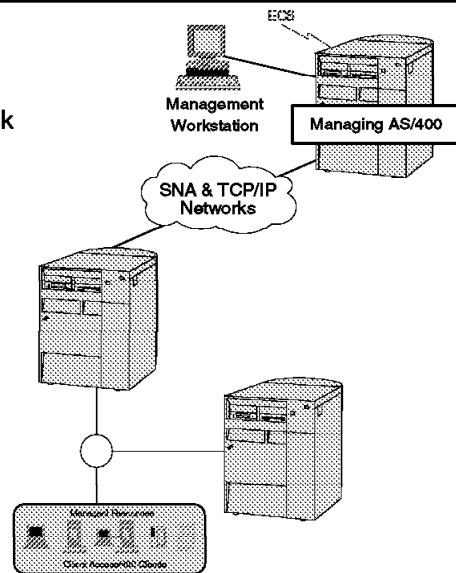
Figure 37. Operational Assistant Display

9.11 OS/400 Alert Support

Alert Support is a problem management product useful for customers who have a network of AS/400 systems and other systems.



- * Detect and log AS/400 system or user problem
- * Notify operator, administrator, help desk
 - Message
 - Pager
- * Analyze problems
- * Send messages, alerts, or traps
- * Filter and direct alerts
- * Check if known problem, fix available
- * Report problem to central site
 - Notify IBM using ECS
- * Receive and apply fix
- * Centralized enterprise problem management and fix distribution
- * Automatic tracking



AACSVL.PRE-22

IBM Corporation 1996

08/96

Figure 38. OS/400 Alert Management

OS/400 Alert Support is useful when:

- You have technical people at only one location in a network.
- You run your own application on your system.

Alert support lets you define your own alert messages so that your applications have the same error-reporting capabilities as the system functions.

- You manage a network with either homogeneous or heterogeneous systems.

Because alerts are designed to be independent of the system architecture, alerts from one system are readable on other systems.

- You must monitor your network status.

Alerts support information about specific network problems that help you track and monitor your system.

- You have unattended remote systems.

Alerts can notify a central site about a problem on an unattended system.

Alerts are generated by OS/400, OS/2, and MVS/ESA. A central AS/400 system can act as the focal point for the alerts allowing the problems to be recognized, logged, and tracked centrally. In addition, with alert filtering, you can direct alerts to the most appropriate support area for quick resolution. With V3R2, V3R7, and subsequent releases you can convert these alerts into SNMP Traps, which can be sent to NetView for AIX.

Alerts and alert filtering are part of OS/400.

A system in your network is designated as the **Service Provider**. This means the other systems in the network can request service in the form of PTFs, for example, from the service provider system. You can also set up *IBMSRV (RETAIN) to make IBM the service provider using ECS (electronic customer support).

9.12 Automating Security Management

Increasing availability on an AS/400 system should also include a secure operating environment. Accidental deletions of files or objects and mischief or sabotage can create an outage just as a system outage or DASD failure can. OS/400 addresses this issue by providing not only integrated security, but also integrated security tools.

OS/400 Security Tools provides security reporting, implementation and management. Enter the GO SECTOOLS command for the following display:

```
SECTOOLS                               Security Tools                               System:  SYSTEMXX
Select one of the following:
  Work with profiles
  1. Analyze default passwords
  2. Display active profile list
  3. Change active profile list
  4. Analyze profile activity
  5. Display activation schedule
  6. Change activation schedule entry
  7. Display expiration schedule
  8. Change expiration schedule entry
                                                                                   More...
Selection or command
===> _____
F1=Help  F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel
```

Figure 39. The First Display of the Security Tools Display

Features of Security Tools include:

- An automated enabling or disabling of user profiles to control access during off hours
- An automatic disabling of user profiles based on company policy
- Monitoring and auditing passwords for company guidelines
- Produce reports showing security violations
- Automation and management of an audit journal

For more information, review the AS/400 publication *AS/400 Tips and Tools for Securing Your AS/400*, SC41-5300-01.

9.12.1 Audit Journal

OS/400 provides an audit journal to log security events occurring on your system. These events are recorded in journal receivers. Different types of security events can be logged into the audit journal, such as the change of a system value or user profile, or an unsuccessful attempt to access an object.

Information in audit journals is used to:

- Detect attempted security violations
- Plan migration to a higher security level
- Monitor the use of sensitive objects, such as confidential files

The audit journal can be set up by using a menu, CL commands, or OS/400 Security Tools. For more information, refer to the *AS/400 Security Reference Guide*, SC41-5301, Chapter 92 “Using the Security Audit Journal.”

For additional security considerations, refer to *Security—Basic V4R1*, SC41-5301, and *AS/400 Tips and Tools for Securing Your AS/400*, SC41-5300-01.

Chapter 10. Work Management for System Availability

The availability of your AS/400 system partly depends on how your work environment is managed. For example, let us assume that you run end-of-day processing every night. If the processing locks the production database for the duration of the job's run time, file changes are impossible. If this end-of-day job does not access all the resources it needs (for example, main storage), processing takes longer. And, if the job has not released the database before the users return to the office, you can experience an unplanned unavailability of the applications that the users require to perform their (human) work.

Proper management of the system's work environment minimize these kinds of outages. This chapter discusses some of the following areas to consider for a higher level of system availability:

- Work with active jobs
- Work control block table
- Display job tables
- System jobs affecting availability
- System values affecting availability
- QSYSOPR message queue wrap when full
- When CPM or dump processing hangs
- When the system date is reset
- End job abnormal
- Reclaim storage
- Other reclaim processes
- Power down system
- Work management APIs

10.1 Work With Active Jobs

The Work With Active Jobs (WRKACTJOB) command allows you to work with both performance and status information for active jobs in the system. User selection on the prompted display changes selection values, such as subsystem (SBS), CPU percent limit (CPUPCTLMT), response time limit (RSPLMT), and automatic refresh interval (INTERVAL). These parameters are shown on the WRKACTJOB prompted command display in Figure 40 on page 140.

The automatic refresh interval parameter specifies the amount of time (in seconds) to wait for an automatic refresh of the display. The default time is 300 seconds (five minutes). Valid values range from 5 to 999 seconds. When pressing the Page Up or Page Down keys, the refreshed values are displayed when the time interval is reached.

When automatic refresh starts, the display is refreshed automatically based on the time specified. The refresh function works similarly to the same function on the WRKSYSACT command in Performance Tools/400. If the user changes the automatic refresh interval value, the value is saved and used as the default for that user.

```

                                Work with Active Jobs (WRKACTJOB)

Type choices, press Enter.

Output . . . . . OUTPUT          *_____

                                Additional Parameters

Reset status statistics . . . . RESET          *NO_____
Subsystem . . . . . SBS           *ALL_____
                                + for more values
CPU percent limit . . . . . CPUPCLMT        *NONE_____
Response time limit . . . . . RSPLMT        *NONE_____
Sequence . . . . . SEQ             *SBS_____
Job name . . . . . JOB            *ALL_____
Automatic refresh interval . . . INTERVAL    *PRV_____

```

Figure 40. New Selection Values for WRKACTJOB

You can change the sequence of the information shown when you place the cursor in the desired column and press the F16 key. You can also choose the Automatic refresh interval option on the WRKACTJOB command, as shown in Figure 40.

This sequencing allows for better system administration when the component of most concern is displayed or listed at the top. For example, when the cursor is in one of the first two positions of the Subsystem/Job column, the display output is ordered by *subsystem*. When the cursor is in one of the last ten positions of the Subsystem/Job field, the display is ordered by *job*.

IBM introduced threads to the operating system in V4R2. A thread is a unit of dispatchable work, each with its own execution environment, such as a call stack. One of the Work with Active Jobs displays shows a column with the thread count. To view this display with the thread information, press F11 twice from the first Work with Active Jobs display. Jobs can be ordered by the number of active threads. The jobs with the largest number of active threads are presented first. Refer to Figure 41 on page 141.

```

                                Work with Active Jobs
                                SYSTEMXX
                                02/08/98 16:04:56
CPU %:   2.3   Elapsed time: 00:01:09   Active jobs: 214

Type options, press Enter.
 9=Exclude 10=Display call stack 11=Work with locks
12=Work with threads 14=Work with mutexes ...

Opt Subsystem/Job  User      Number  Type  CPU %  Threads
---
--- QBATCH        QSYS      031412  SBS   .0     1
--- NEWLOOK       A970107A  031816  BCH   .0     1
--- PRT_ORDER     A970107A  031638  BCH   .0     1
--- TTLMOM        XXTOOL    031920  BCH   .0     1
--- QCMN          QSYS      031413  SBS   .0     1
--- QCTL          QSYS      031375  SBS   .0     1
--- QJSCCPY       A97011359 032004  BCH   .0     1
--- QSYSSCD       QPGMR     031402  BCH   .0     1
--- QINTER        QSYS      031408  SBS   .0     1
--- QPADEV0007    DMT       031779  INT   .0     1
--- QPADEV0009    A970107A  031967  INT   .0     1
--- QPADEV0016    A97011359 032002  INT   .8     1
--- QSERVER       QSYS      031396  SBS   .0     1
--- QPWFSERVSD    QUSER     031549  BCH   .0     1
--- QPWFSERVS0    QUSER     031410  PJ    .0     1
--- QPWFSERVS0    QUSER     031561  PJ    .0     1

More . . .

===>
F19=Start automatic refresh  F21=Nondisplay instructions/key
F24=More keys

```

Figure 41. Work with Active Jobs

Thread identifiers are also visible when the call stack is displayed for a job. To see a job's thread information, select option 20 from the Work with Job (WRKJOB) menu to *Work with threads (if active)*. Or select option 12 to *Work with threads* from the WRKACTJOB screen.

With the inclusion of thread information, there are three additional wait indicators:

- CNDW—Waiting on a handle based condition
- HLDT—The initial thread of the job is held
- THDW—The initial thread is waiting for another thread to complete

10.2 Work Control Block Table

The Work Control Block Table (WCBT) is a system object that maintains information for all jobs from the time they enter the system until they are removed. Each job is represented by a single entry in the WCBT.

On V4R1 and later systems, an index is maintained over the WCBT. This allows the system to locate entries much faster than on releases prior to V4R1 that use a sequential search algorithm.

The WCBT consists of a single space that contains the header information and one to ten spaces of entries for jobs. The WCBT can contain a maximum of

163520 entries, which means that the maximum number of jobs on an AS/400 system is limited to 163 520 on V4R2 systems. There is a Work Control Block Table Entry (WCBTE) for every job tracked by OS/400. Jobs are counted if they are:

- On the job queue ready to run
- Actively in the system doing work
- Completed with output that remains on an output queue

When a job is no longer tracked by the system, a new job that is starting can reuse its Work Control Block Table Entry (WCBTE). However, there can be situations on heavily utilized systems where the number of jobs being tracked is quite large. This causes the total storage occupied by the WCBTEs to increase.

If a large number of these entries (jobs) become “no longer tracked,” it can take a while to reuse the “available” entry space. During that time, processing the Work Control Block Table entries with commands, such as Work With Subsystem Jobs (WRKSBSJOB) or the List Job (QUSLJOB) API, can become time consuming on systems with a large number of tracked jobs. Eventually, you need to compress these entries to remove the empty slots to return to the count of entries specified with the QTOTJOB system value. The compression is performed during IPL. Compression is requested using the compress job tables (CPRJOBTL) parameter on the Change IPL Attributes (CHGIPLA) command shown in Figure 42. Compression reduces the size of the WCBT by freeing up the unused WCBTEs.

```

                                Change IPL Attributes (CHGIPLA)

Type choices, press Enter.

Restart type . . . . . RESTART          *SYS
Keylock position . . . . . KEYLCKPOS     *NORMAL
Hardware diagnostics . . . . . HDWDIAG   *MIN
Compress job tables . . . . . CPRJOBTL   *NONE
Check job tables . . . . . CHKJOBTL      *ABNORMAL
Rebuild product directory . . . RBDPRDDIR *NONE
Mail Server Framework recovery MSFRCY    *NONE
Clear job queues . . . . . CLRJOBQ       *NO
Clear output queues . . . . . CLROUTQ    *NO
Clear incomplete joblogs . . . CLRINCJOB  *NO
Start print writers . . . . . STRPRTWTR  *YES
Start to restricted state . . . STRRSTD    *NO

```

Figure 42. Change IPL Attributes (CHGIPLA) Display

When the system starts using the tenth Work Control Block Table, a CPI1468 message is sent indicating that the WCBT is nearing capacity. The following recovery operation is recommended in the message text:

```

. . . If the work control block table is permitted
to fill completely, then jobs cannot be successfully submitted and
the subsequent IPL may result in failure with resulting loss of
all spooled files and jobs on job queues. Delete unneeded spooled
files or end unneeded jobs on job queues to free up job structures...

```


For further information, refer to *Work Management*, SC41-5306.

10.2.1 Work Control Block Table Cleanup

QWCBTCLNUP is a cleanup system job that runs during the IPL for Work Control Block Table entry cleanup. This system job usually completes processing prior to the end of the IPL, but it can run after the IPL on systems where more cleanup is required.

Work Control Block Table entries are also compressed by using the CHGIPLA command described in Section 4.5, “Changing IPL Attributes” on page 42.

10.2.1.1 Compress Job Tables at IPL for Systems Prior to V4R1

The code to compress job tables is in the base operating system for systems after V3R1. To provide the same function as the CPRJOBTL parameter on the CHGIPLA command, order and apply Programming Temporary Fix (PTF) SF23320 for V3R1 systems to compress the job tables.

The original V3R1 PTF has been replaced several times. However, the only user documentation on how to compress the job table at IPL is in the cover letter for SF23320. We recommend that you print the cover letter for PTF SF23320. The cover letter describes how to activate this function as follows:

```
CRTDTAARA DTAARA(QSYS/QWCBTCMPTB) TYPE(*CHAR) LEN(1)
```

```
CHGDTAARA DTAARA(QSYS/QWCBTCMPTB) VALUE('1') /* COMPRESS ON */
```

```
CHGDTAARA DTAARA(QSYS/QWCBTCMPTB) VALUE('0') /* COMPRESS OFF */
```

For instructions on printing a cover letter, refer to *System Operation*, SC41-4203.

10.3 Display Job Tables

The Display Job Tables (DSPJOBTL) command is a useful tool to monitor the growth of structures so that an adjustment to WCBT related system values can be made proactively. This command helps prevent scheduled or unscheduled outages caused when the WCBT becomes full. DSPJOBTL provides an interface to Work Control Block Table entries in the WCBTs.

The number of entries in these tables can affect the performance of various IPL steps, OS/400 commands, and Application Program Interfaces (APIs) that work with jobs. If your AS/400 system seems to hang for no apparent reason, check the status of the WCBT settings.

To check the status on systems prior to V4R1:

- Obtain the number of jobs in the system from the WRKSYSSTS display.
and
- Obtain the number of active jobs in the system from the WRKACTJOB display.
and
- Compare these values to the following WCBT related system values to see whether any of the system-value user-defined limits are exceeded.
 - QTOTJOB
 - QADLTOTJ

- QACTJOB
- QADLACTJ

On V4R2 systems, you can accomplish all of this by entering the following command:

DSPJOBTL

Unlike the WRKACTJOB and WRKSYSSTS displays, the Display Job Tables display shows information from internal system objects and presents a clearer picture of job structures (see Figure 43).

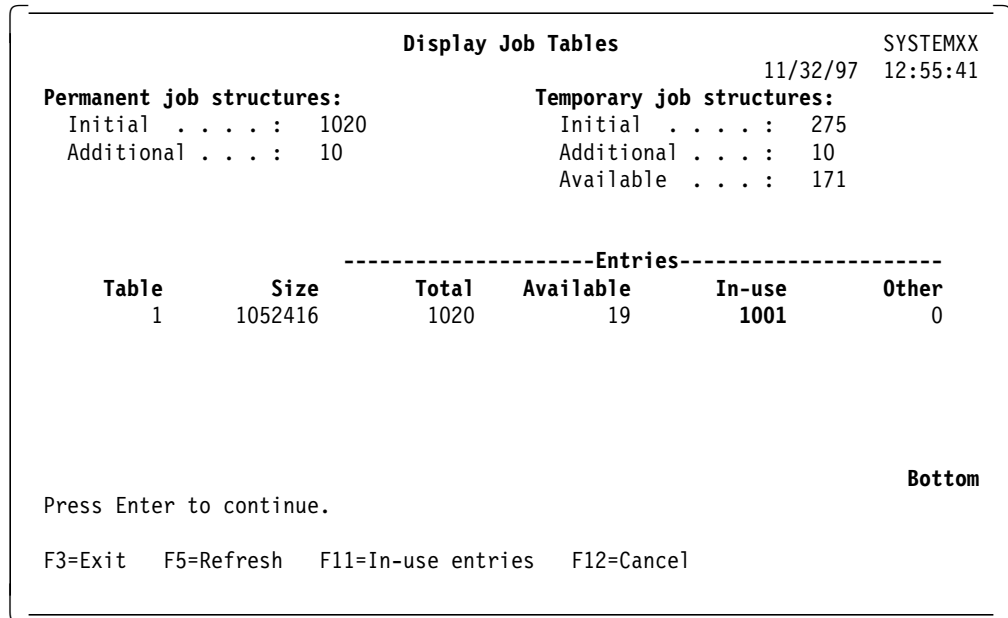


Figure 43. Display Job Tables

The permanent job structures, temporary job structures, and entries sections of the DSPJOBTL output are described in the following sections.

10.3.1 Permanent Job Structures Field

A permanent job structure is assigned to a job when it enters the system. The permanent job structure is not available for reuse until the job ends either without spooled output or all of the spooled output for a job is printed. The **initial** and **additional** values are displayed.

Initial This value represents the initial number of entries in the job tables for which the auxiliary storage is allocated during an IPL. This value displays the current setting of the QTOTJOB system value.

Additional This value indicates the additional number of entries for which auxiliary storage is allocated when there are no more available entries in the job tables. Although you can set this value greater than 500, a maximum of 500 job table entries are added at one time. The additional value displays the current setting of the QADLTOTJ system value.

10.3.2 Temporary Job Structures Field

A temporary job structure is assigned to a job when it becomes active. When the job ends, the next job that becomes active can reuse the temporary job structure. This storage is in addition to the storage allocated when the system value QTOTJOB is used. The **initial**, **additional**, and **available** values are displayed.

- Initial** This value is the initial number of temporary job structures for which storage is allocated during an IPL. It indicates the current setting of the QACTJOB system value.
- Additional** This value is the additional number of temporary job structures for which storage is allocated when all available temporary job structures are assigned to active jobs. It displays the current setting of the QADLACTJ system value.
- Available** This value represents the number of temporary job structures that are created, but not assigned to active jobs. These jobs structures are available for use when jobs become active.

10.3.3 Entries Field

When you press the F11 key on the Display Job Tables (DSPJOB_TBL) display, information appears about the in-use entries as shown in Figure 44. Note that you can determine how many entries are consumed by active jobs, job queue jobs, and output queue jobs.

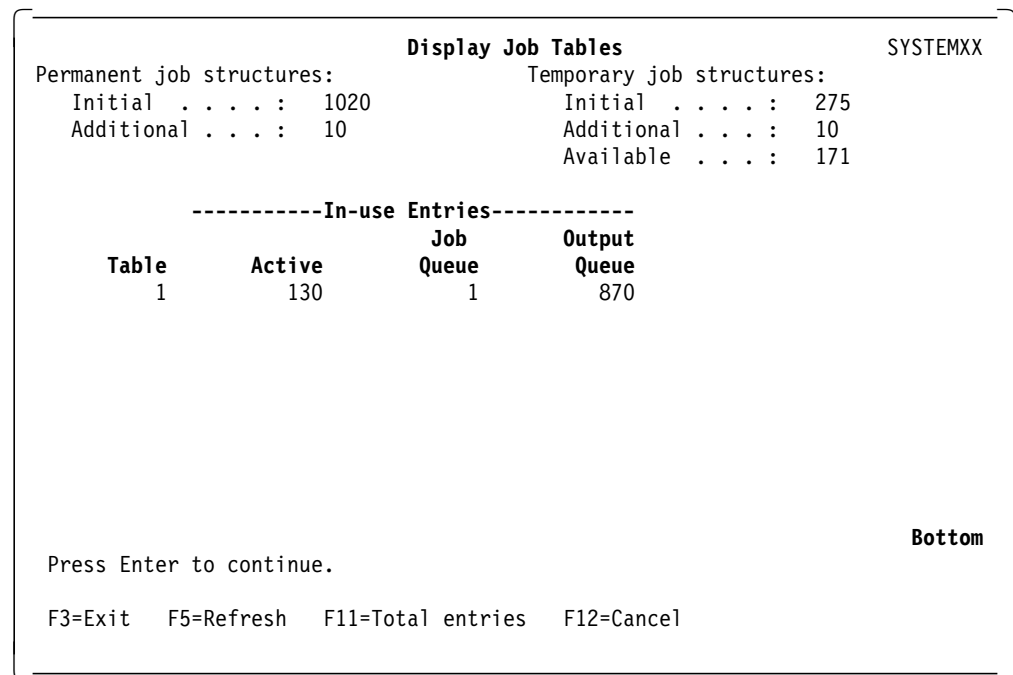


Figure 44. In-Use Entries on Display Job Tables Display

The *In-use entries* field provides information about the entries contained in the job tables (WCBTs).

- Table** This value shows the number of the WCBT. The values range from one to ten since there are a maximum of ten WCBTs per system.
- Size** This value indicates the size of the WCBT in bytes.

- Total** This value indicates the total number of entries contained in the WCBT. If the number of available entries is large, you can compress the job tables during the next IPL. Use the Change IPL Attributes (CHGIPLA) command as described in Section 4.5, "Changing IPL Attributes" on page 42, to change the option used to compress the job tables.
- Available** This value indicates the number of entries that are available for new jobs. If there are *no* available entries, the system can experience a *severe performance degradation* when it starts a new job because the table has to be extended. However, too many available entries degrade the performance of IPL steps that process the WCBT and run-time functions that work with jobs.
- In-use** This value shows the number of entries currently used by jobs that are on a job queue, jobs that are active, or jobs that have completed but spooled output on an output queue. This total is calculated as shown here:

A Simple Calculation of In-Use Entries	
130	Active +
1	in Job Queue +
870	in Output queue =
1001	The Number of In-Use Entries

Notice that the number of in-use entries total matches that shown in Figure 43 on page 144.

- Other** This value shows the number of entries that are not available and not currently in use by jobs that are on a job queue, jobs that are active, or jobs that have completed but spooled output on an output queue. It includes entries that are marked as unusable and those that are for jobs in transition (such as from a job queue to the active state).

10.4 System Jobs Affecting Availability

Many of the system jobs that relate to availability and recovery are highlighted on the WRKACTJOB display in Figure 45 on page 147.

Work with Active Jobs							SYSTEMXX	
						12/16/97	12:46:45	
CPU %:	1.5	Elapsed time:	02:45:24	Active jobs:	228			
Opt	Subsystem/Job	User	Type	CPU %	Function	Status		
—	QALERT	QSYS	SYS	.0		DEQW		
—	QCMNARB01	QSYS	SYS	.0		EVTW		
—	QCMNARB02	QSYS	SYS	.0		EVTW		
—	QDBSRVXR	QSYS	SYS	.0		DEQW		
—	QDBSRVXR2	QSYS	SYS	.0		DEQW		
—	QDBSRV01	QSYS	SYS	.0		EVTW		
—	QDBSRV02	QSYS	SYS	.0		DEQW		
—	QDBSRV03	QSYS	SYS	.0		DEQW		
—	QDBSRV04	QSYS	SYS	.0		DEQW		
—	QDBSRV05	QSYS	SYS	.0		DEQW		
—	QDCPOBJ1	QSYS	SYS	.0		EVTW		
—	QDCPOBJ2	QSYS	SYS	.0		EVTW		
—	QFILESYS1	QSYS	SYS	.0		TIMW		
—	QJOBSCD	QSYS	SYS	.0		EVTW		
—	QLUR	QSYS	SYS	.0		EVTW		
—	QLUS	QSYS	SYS	.0		EVTW		
—	QPFRAJ	QSYS	SYS	.0		EVTW		
—	QQQTEMP1	QSYS	SYS	.0		DEQW		
—	QQQTEMP2	QSYS	SYS	.0		DEQW		
—	QSPLMAINT	QSYS	SYS	.0		EVTW		
—	QSYSARB	QSYS	SYS	.0		EVTW		
—	QSYSARB2	QSYS	SYS	.0		EVTW		
—	QSYSARB3	QSYS	SYS	.0		EVTW		
—	QSYSARB4	QSYS	SYS	.0		EVTW		
—	QSYSARB5	QSYS	SYS	.0		EVTW		
—	QSYSCOMM1	QSYS	SYS	.0		EVTW		
—	Q400FILSVR	QSYS	SYS	.0		DEQW		
—	SCPF	QSYS	SYS	.0		EVTW		
	.							
	.							
	.							
	.							
====>							Bottom	
F21=Display instructions/keys								

Figure 45. A Partial List of System Jobs

This section describes:

- QSYSARB and QSYSARBN—System Arbiters
- QJOBSCD—Job Scheduler

QCMNARBn, QPFRAJ, QSYSCOMM1, and QDBSRVnn jobs are described throughout other sections of this redbook.

For more details about the other system jobs, refer to the *Work Management Guide*, SC41-5306.

10.4.1 QSYSARB and QSYSARBN Jobs

The SCPF system job starts the system arbiter jobs, which remain active until the system ends. The system arbiters provide the environment for running high-priority functions. They allow subsystems to start and end and keep track of the state of the system (for example, whether the system is in a restricted state).

The system arbiters are identified by job names QSYSARB and QSYSARB2 through QSYSARB5. They are central to error handling. They are the highest priority jobs within the operating system. Each system arbiter responds to system-wide events that must be handled immediately and to those that can be handled more efficiently by a single job rather than multiple jobs.

QSYSARB is also responsible for starting the Logical Unit Services (QLUS) job during an IPL.

10.4.2 QJOBSCD—Job Scheduler

The job scheduler system job QJOBSCD controls the job scheduling functions of the system. QJOBSCD monitors the timers for job schedule entries and scheduled jobs. It starts during IPL and remains active until the system is powered down.

QJOBSCD uses the job-schedule object QDFTJOBSCD. QDFTJOBSCD is stored in the QUSRSYS library with an object type of *JOBSCD. The job-schedule object contains entries that set up a schedule of jobs. You schedule jobs by adding a job-schedule object. You cannot create, delete, rename, or duplicate the job-schedule object. And, you cannot move it to another library.

The job-schedule object is saved with the Save Library (SAVLIB), Save Object (SAVOBJ), or Save Changed Object (SAVCHGOBJ) commands. You can restore it with the Restore Library (RSTLIB) or Restore Object (RSTOBJ) command. Restoring the job-schedule object updates the next submission date for each entry. You can restore the job-schedule object to the system from which it was saved or to a different system, but you cannot restore it to a library other than QUSRSYS. If you restore the job-schedule object to a different system, the job submission history is cleared in each entry.

If there are scheduled jobs that are not run when they should be run, the administrator must make sure that the QJOBSCD job is still active. If QJOBSCD does not appear as a system job on the WRKACTJOB display, you can only start it with an IPL. You can release jobs from the job schedule manually until an IPL is scheduled.

Note: IBM has a Job Scheduler/400 product that further enhances job scheduling options on the AS/400 system. Refer to *Job Scheduler for OS/400*, SC41-4324, for more information about automating operations using IBM's job scheduler.

10.5 System Values Affecting Availability

If auxiliary storage exceeds threshold values, availability of the system and network functions become unpredictable.

Note: Exceeding other thresholds can also affect system availability. Refer to Appendix A, "AS/400 Maximum Capacities" on page 361 for further information.

The use of two system values helps to monitor and manage the growth of auxiliary storage:

- Auxiliary Storage Lower Limit—QSTGLOWLMT
- Auxiliary Storage Lower Limit Action—QSTGLOWACN

These system values are described later in this section, as well as device recovery for local devices. See Section 12.2.3, “Device Recovery Performance for Display Devices” on page 185, to learn more about communications device recovery considerations.

Note: The QPWRDWNLMT system value also affects availability and recovery. We discuss this topic separately in Section 10.12.3, “Timeout Options for Power Down System” on page 162.

10.5.1 Auxiliary Storage Lower Limit—QSTGLOWLMT

The percent of storage currently used in the system ASP appears on the Work with System Status (WRKSYSSTS) display shown in Figure 46.

Work with System Status										SYSTEMXX
										12/16/97 24:31:34
% CPU used :	2.4	System ASP :	20.45 G							
Elapsed time :	00:00:01	% system ASP used :	86.4620							
Jobs in system :	463	Total aux stg :	20.45 G							
% perm addresses :	.007	Current unprotect used :	1552 M							
% temp addresses :	.010	Maximum unprotect :	1621 M							
Sys Pool	Pool Size K	Reserved Size K	Max Act	----DB----- Fault Pages	--Non-DB--- Fault Pages	Act- Wait	Wait Inel	Act- Inel		
1	<u>61304</u>	35376	+++++	.0 .0	8.5 8.5	.0	.0	.0		
2	<u>399952</u>	0	<u>18</u>	.0 .0	.0 .0	139.8	.0	.0		
3	<u>5240</u>	0	<u>5</u>	.0 .0	.0 .0	.0	.0	.0		
4	<u>57792</u>	28	<u>36</u>	.0 .0	5.4 6.2	46.6	.0	.0		
										Bottom
===>										
F21=Select assistance level										

Figure 46. Work with System Status

On V4R2 systems, the administrator can use the QSTGLOWACN and QSTGLOWLMT systems values to determine what action the system takes when a minimum percent of storage is reached. The system performs the action when the threshold is reached and messages are sent to the operator. If the low storage situation is not rectified, some system functions, such as SNADS, halt processing. Eventually the system stops.

When the lower limit specified with the auxiliary storage lower limit (QSTGLOWLMT) system value is reached, an action is taken. The auxiliary storage lower limit action (QSTGLOWACN) system value specifies the action associated with this limit.

Note: QSTLOWLMT does not affect *user* ASP limits.

10.5.2 Auxiliary Storage Lower Limit Action

In releases prior to V4R2, the system terminates abnormally when the ASP storage capacity is completely filled. On V4R2 systems, users can decide what actions the AS/400 system performs when the remaining storage capacity specified by the QSTLOWLMT system value is reached. The actions that you can specify with the auxiliary storage lower limit action (QSTGLOWACN) system value include:

- *MSG** Sends message CPI099C to the QSYSMSG and QSYSOPR message queue to indicate *Critical storage lower limit reached*.
- *CRITMSG** Sends critical message CPI099B to indicate *Critical storage condition exists* to the user specified in the service attribute to receive critical messages.
- *REGFAC** Submits a job to call exit programs registered for the QIBM_QWC_QSTGLOWACN exit point.
- *ENDSYS** Ends the system to a restricted state.
- *PWRDOWNSYS** Powers down the system immediately and restarts it.

Monitoring these two system values provides the administrator more flexibility for managing system availability relative to storage capacity.

The WRKASP utility is a tool useful to monitor and manage user ASPs. Refer to Section 9.4, "WRKASP Utility" on page 120, for information on WRKASP.

10.5.3 Device Recovery Action

Various situations can cause a workstation to lose connection with the AS/400 system. How and how soon a workstation resumes availability and what diagnostics are available for problem determination depend on the device I/O error action (QDEVRCYACN) system value. They also depend on:

- What problem determination tools are available
- How recovery was planned
- How soon the problem was noticed

The QDEVRCYACN system value allows several options to handle the disruption in an interactive job workstation. The possible actions are:

- *DSCMSG** Disconnects the job. When signing on again, an error message is sent to the user's application program. The default value is *DSCMSG.
- *MSG** Signals the I/O error to the user's application program. The application program performs error recovery.
- *DSCENDRQS** Disconnects the job. When signing on again, a cancel request function is performed to return control of the job back to the last request level.
- *ENDJOB** Ends the job. A job log is produced. A message is sent to the job log and QHST log. Job priority is lowered by ten, the timeslice is set to 100 milliseconds, and the purge attribute is set to *YES. This operation minimizes the performance impact of ending the job.

***ENDJOBNO LIST** Ends the job and, therefore, does not produce a job log for it. A message is sent to the QHST log.

Recommendation

The system administrator should review the current values of the QDEVRCYACN system value and consider its implication. Be aware that disconnected jobs consume resources.

Recovery for communications devices is described in Chapter 12, “Communications Error Recovery and Availability” on page 181, of this redbook.

10.6 QSYSOPR Message Queue Wrap When Full

The QSYSOPR message queue is a main source of information for the system operator to investigate problems on the AS/400 system. As a message queue, however, when the size limit is reached, it can become full and the job can loop. Recovery can require an IPL.

Message queues for jobs can wrap when they fill if the QJOBMSGQFL job message queue full action system value is set to *WRAP or *PRTWRAP. This system value, however, does *not* affect QSYSOPR. AS/400 system administrators and system operators avoid a full condition on QSYSOPR by carefully monitoring the queue, routinely clearing messages that are not answered, and monitoring for a message indicating the queue is reaching a full condition.

With the application of a PTF, it is now possible to allow QSYSOPR to wrap. The PTF number is identified in Table 10.

Release	PTF	APAR
V4R2	SF45613	SA68461
V4R1	SF44163	SA68461

These PTFs allow QSYSOPR to wrap if the QSYSOPRFUL data area in library QUSRSYS exists. If the data area exists, QSYSOPR wraps by writing over the oldest informational and answered messages on the queue. If more room is required, the oldest unanswered inquiry messages are responded to with their default replies then they are removed from the queue. If QSYSOPR wraps and there still is no room for the another message, a wrap message is sent to QHST and to the job log of the job sending the message that does not “fit” in QSYSOPR.

To end the wrapping of QSYSOPR, the data area can be deleted. If the data area does not exist and QSYSOPR fills, the system responds as without the PTF. That is, a CPF2460 escape message indicating “message queue &1 could not be extended” is issued to the program or user sending a message to the full queue.

Refer to the PTF cover letter for special considerations when this PTF is installed, including what happens when *PUBLIC authority to the QCPFMSG message queue is set to *USE.

10.7 When CPM or Dump Processing Hang

When CPM or main store dump processing hangs or otherwise fails, you can select F34 to force a retry of CPM or MSD processing. Known as an MSD or CPM IPL, F34 performs an IPL again on the system while preserving the contents of the main storage. Flags are set to automatically retry the CPM or main store dump process. The dump is useful for problem determination because with many hang situations, it is the only tool available for problem analysis.

Recommendation

Use F34 instead of F08 to perform a shutdown. F08 is a forced power down and no dump is produced for additional problem analysis.

You can also access main storage dump information so you can view or save it before storage management recovery. An option to automatically copy the main storage dump to the system ASP and IPL is available for V4R1 systems and later. Using Dedicated System Tools or System Service Tools, select the Main Storage Dump Manager to enable the autocopy option.

10.8 When the System Date is Reset

Periodically, the system date is reset through means other than operator intervention, such as:

- When an MFIOP is replaced or unplugged
- When the system control panel is removed or unplugged

In these situations, the system date is set to a default date of August 23, 1928. This causes:

- Scheduled jobs to not run
- Applications to produce erroneous results
- User profiles with expired passwords

On V4R2 systems, some maintenance activities occur that cause the system to set the date to its default. When this happens, the system is forced into an attended mode IPL where the operator must enter a valid date for the IPL to continue. A message is sent to QHST and QSYSOPR message queues to indicate what happened in the IPL.

10.9 End Job Abnormal

Periodically, jobs are in such a state that even an End Job Abnormal (ENDJOBABN) command does not force the job to end. This condition indicates a software defect, and you must report it to the IBM AS/400 Support Line.

Apply the following PTFs to your system to ensure that the enhancements to the ENDJOBABN process are included. These enhancements offer additional code to force more jobs to end than was previously possible. Refer to the associated APAR for more information and apply considerations.

<i>Table 11. PTFs and APARs to Improve ENDJOBABN</i>		
Release	PTF	APAR
V4R2	SF45302	SA69691
V4R1	SF45312	SA69691
V3R7	SF45256	SA69627
V3R2	SF45301	SA69690

Note: The PTFs listed contain the code when the enhancements were initially available. When ordering the PTF, any PTF which superseded it is sent in its place automatically. Also note that there is no fix for R360.

If an ENDJOBABN fails to end a job when the PTFs are applied, the proper documentation is available for additional problem determination by IBM AS/400 Support Line representatives and defect support.

10.10 Reclaim Storage

Reclaim storage is a maintenance routine that you must run on an AS/400 system from time to time. The Reclaim Storage (RCLSTG) command is used to:

1. Perform a cleanup when database cross references are in error.
2. Correct abnormal conditions for objects in auxiliary storage that may be affected by an unexpected termination of hardware or programs (for example, a failed restore or installation activity).
3. Correct, where possible, objects that are incompletely updated because of an unexpected termination.
4. Resolve user profiles containing incorrectly recorded object ownership (for example, if objects are on the system but their owning user profile or authorization list is damaged or lost on the system). This symptom is apparent when you use the WRKOBJOWN command and the Objects By Owner display does not have an associated library.
5. Delete any unusable (damaged) objects or fragments. For database files, this symptom is reported by a CPF8113 message that indicates damage to a database file member.
6. Perform a general cleanup of auxiliary storage after a failed restore or installation activity.

Run RCLSTG for system maintenance to help regain storage space. If the percent of auxiliary storage space used seems unexplainably high, schedule time to run a RCLSTG. Since it requires a dedicated system to perform RCLSTG functions and takes a long time to process, this process is difficult to schedule.

Note: Report any Reclaim Storage process that takes longer than 12 hours to complete to the IBM Software Service (AS/400 Support Line) for analysis of program defects.

For damage to the database cross reference files, the system directs the user to perform the storage reclaim. This assumes that the user follows the recovery action described in the error messages that alert the operator to this problem. These messages include:

- CPF32A1—*System cross reference queue deleted and created again*

- CPF32A2—*Unexpected data found in system cross reference files*
- CPF32A3—*System file &1 in library &2 cannot be opened*
- CPF32A4—*Internal failure in system cross reference program*

The system operator can also use the RCLSTG command to reclaim storage when enough storage is not available during an IPL to make the system fully operational. The system operator can specify the command immediately after receiving the message about insufficient storage.

Only permanent objects in auxiliary storage are reclaimed with the RCLSTG command. Temporary objects are reclaimed by performing an IPL. If little additional auxiliary storage is available, the system overhead required to run a Reclaim Storage command may require more than the remaining storage. This causes RCLSTG to fail.

10.10.1 The Benefits of Running RCLSTG

When an interactive job is complete, the QTEMP library of the job is not deleted immediately, but cleared during the next IPL. Ending a job with ENDJOB OPTION(*IMMED) causes an abrupt job termination that usually leaves some temporary objects as orphans (without a library or owner). These orphans are cleared during the execution of RCLSTG.

The same principle applies to abnormal system termination, only on a larger scale. Besides all the user jobs, when the system jobs are terminated abnormally, a multiple amount of orphans occur.

To help you understand RCLSTG, review the types of actions that the reclaim function performs and the corrective actions that it takes when possible:

1. For each object on the system, RCLSTG:
 - Ensures that it is owned by a profile.
 - Ensures that it is in a library.
 - Ensures that it is not secured by a damaged or destroyed authorization list, and assigns the objects to the system supplied authorization list—QRCLAUTL.
 - Deletes objects that are not completely created.
 - Validates object storage counts in all user profiles.
 - Deletes duplicate user profiles.
 - Deletes subsystem descriptions that are incomplete or with header damage.
 - Recovers the ASP.
 - Deletes duplicate configuration descriptions.
 - Deletes internal system objects if they are not needed.
 - Deletes invalid objects.
2. For each library related object on the system, RCLSTG:
 - Checks it for damage.
 - Ensures all pieces of a library are connected.
 - Removes object description entries if there is no corresponding entry in the library.
 - Deletes any partially created library.
 - Deletes duplicate libraries.
3. For each office related object on the system, RCLSTG:
 - Ensures DLO objects are in QDOC.

- Ensures DLO objects are owned.
 - Deletes DLO objects that are invalid.
 - Verifies the mail logs on the system.
 - Ensures all mail is on someone's mail log.
 - Ensures all mail documents are attached to a mail item.
 - Performs functions different from RCLDLO.
4. For each database-related object on the system, RCLSTG:
- Performs any pending commit recovery.
 - Deletes any secondary objects that are danglers.
 - Rebuilds database files for which data exists.
 - Rebuilds database cross reference files.
 - Identifies database formats and directories that are orphans (no owner or library).
 - Recovers databases.

Review these items with serious consideration to develop a regular schedule for running RCLSTG.

10.10.2 Reclaim Storage Options

Beginning with V4R1, you can specify a parameter to control which parts of the system to reclaim. The added parameters enable a staging of the RCLSTG function. The two added parameters, as shown in Figure 47, are:

- Select
- Omit

These options allow administrators to plan for multiple, smaller system outages as opposed to a single, longer outage for running the reclaim.

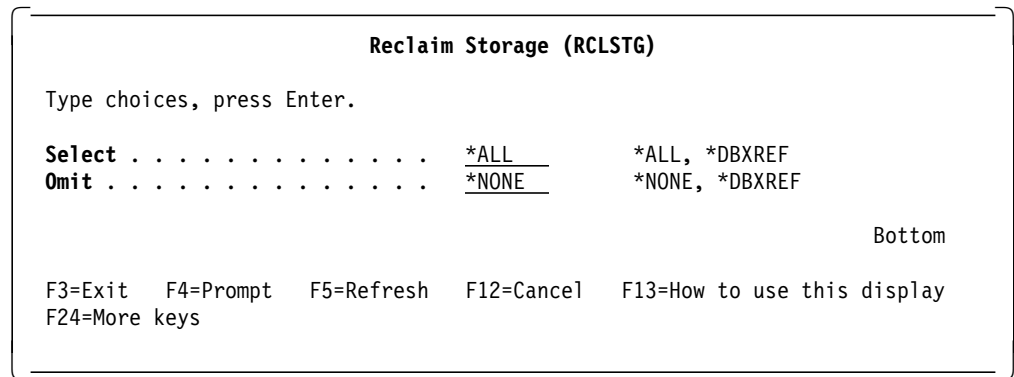


Figure 47. New Parameters for the Reclaim Storage Command

The select parameter with a *DBXREF value specified allows the user to reclaim only the database cross-reference files without performing a complete reclaim function. The time it takes to rebuild these tables depends on the number of database objects on the system.

The omit parameter offers a way to omit the database cross-reference rebuild from a complete reclaim storage operation.

The default values of *SELECT(*ALL) OMIT(*NONE) perform a complete reclaim storage function, such as RCLSTG on systems prior to V4R1. On these systems, to reclaim the database cross-reference tables, run a mini-reclaim storage by entering:

CALL QDBRCLXR

This function only takes a few minutes, but requires a restricted state. It is functionally equivalent to the command:

RCLSTG SELECT(*DBXREF)

If running QDBRCLXR after an IPL, wait several minutes for the IPL to complete. This enables IPL-related cleanup and startup activities to complete normally.

Note: RCLSTG requires restricted state processing regardless of what is reclaimed.

10.10.3 Reclaim Storage Status Messages

The RCLSTG procedure runs a long time because it examines every object on the system. It does not make changes to an object or move it to the QRCL library unless it detects that the object header (the container) is damaged or lost (not in a library).

Because the amount of time required to run this command varies, the system periodically sends messages to the workstation where the command was initiated. The RCLSTG command issues the following sequence of status messages. From this list, you can deduct *how far along* the reclaim process is:

- CPI8220** Message queue QSYSOPR in *HOLD delivery mode.
- CPI8212** Database/library/directory recovery in progress.
- CPI8213** Processing objects on the system.
- CPI8206** &% of objects processed.
- CPI8210** Processing database relationships.
- CPI8218** Directory recovery in progress.
- CPI8219** Directory cleanup in progress.
- CPI8215** Object description verification in progress.
- CPI8217** Mail Server Framework cleanup in progress.
- CPI8216** Final cleanup in progress.
- CPI8214** All permanent objects have valid owners.
- CPC8208** RCLSTG processing complete. &1 objects processed. &2 objects deleted.

Note: The messages listed indicate when a particular step is reached in the process. They do *not* indicate that the reclaim process is one-third complete when you see CPI8206 (message four of twelve).

As you can see from the message descriptions, RCLSTG analyzes database objects, library objects, document library objects, mail logs, and more. It does *not*, however, detect all damage. For example, RCLSTG does not scan the data space or data space indexes (logical files) for damage. The RCLSTG procedure does not detect or correct internal damage to objects, such as record-level damage in database files.

10.10.4 Reclaim Storage Error Messages

If something goes wrong during the reclaim, some of the following error messages may appear:

- CPF2119** Library &1 locked.
- CPF2120** Cannot delete library &1.
- CPF2126** Attempt to recover library &1 failed.
- CPF2127** User profile &2 damaged.
- CPF8201** User profile &1 does not exist or is damaged.
- CPF8204** Commitment control cannot be active during reclaim storage.
- CPF8205** Library &1 does not exist or is damaged.
- CPF8209** System not in proper state to reclaim storage.
- CPF8211** Library &1 damaged. RCLSTG command ended.
- CPF8224** Duplicate object found while moving or renaming member.
- CPF8251** RCLSTG command ended. Library &1 damaged.
- CPF8252** Error occurred during rebuild of damaged library &1.

Each messages carries a security level attribute of 40 or 50 and, therefore, causes a halt in RCLSTG processing.

10.10.5 Completion Time for Reclaim Storage

The best estimate to determine how long a reclaim process takes to complete is derived by running it on your system or see how long it took the last time you ran it. To determine how long the reclaim process took the last time you ran it, view the QRCLSTG data area:

DSPDTAARA DTAARA(QUSRSYS/QRCLSTG)

Compare the start and end date and time. The start and end date are in the first two fields of the data area. You also see the release of the operating system, the system name, and serial number when displaying the data area (see Figure 48).

Data area	:	QRCLSTG
Library	:	QUSRSYS
Type	:	*CHAR
Length	:	200
Text	:	Information from RCLSTG
Offset	Value	
0	'0971102 152519	'
50	'0971102 171907	'
100	'V4R2MO SYSTEM03 0526SXP	'
150	'	'

Figure 48. QRCLSTG Data Area

Based on what happens during a reclaim storage process, consider the following factors, which affect how long it takes RCLSTG to run, to estimate the time to execute on your system:

- Amount of free disk space on the AS/400 system
- Amount of jobs ended with ENDJOB(*IMMED)
- Number of abnormal system terminations
- Number of objects on the system
- Number of damaged objects
- Number of user profiles on the system
- Types of objects on the system
- Amount of damage to relational objects
- Object size
- Amount of main and auxiliary storage
- Percentage of auxiliary storage in use
- Amount of time since the last IPL
- Amount of time since the last RCLSTG command was run

If you begin the reclaim storage process and determine that it will not complete in the time you allotted for it, you can end the function by issuing a System Request. Select option 2 to *End Previous Request*.

Note: Cancelling the RCLSTG procedure is not recommended, but should not cause any loss of data or damage to objects. It can occasionally cause problems with a database file or problems with SQL or IDDU applications if RCLSTG ended while the cross-reference files are rebuilt. However, the risk is considered low.

When RCLSTG is restarted, it starts all over again. The time you may have saved is from any object recovery RCLSTG accomplished before it was cancelled.

If the auxiliary storage usage is high, and you do not have reason to suspect object damage, it is unlikely that a RCLSTG can free up significant disk space.

10.11 Other Reclaim Processes

In addition to RCLSTG, spooled files and document library objects (DLOs) periodically need to be reclaimed. This section discusses these processes.

10.11.1 Reclaim Document Library Object (RCLDLO)

RCLDLO *ALL is a disaster recovery technique for lost DLO related database files in QUSRSYS and user ASPs. With RISC systems allowing up to five million DLOs, the time to run RCLDLO can be significant (three 25K DLOs per hour).

On V4R2 systems, an additional value is offered for the reclaim document library object (DLO) parameter to allow greater efficiency. The *DOCDDL value for DLO specifies that internal document library system objects *and* document details are to be reclaimed. DLO(*DOCDDL) synchronizes the relationships between all internal system objects, document details, and DLOs.

On systems prior to V4R2, a RCLDLO DLO(*ALL) option is taken to perform this processing. The *ALL option involves extra processing that is unnecessary for some recovery scenarios.

10.11.2 Reclaim Spool Storage (RCLSPLSTG)

Spooled files are system objects that reside in the QSPL library. When a spooled file is printed, not held, or deleted from the output queue, the space used by the actual database member is not freed up until a Reclaim Spool Storage (RCLSPLSTG) command is run.

Events that trigger the RCLSPLSTG command are:

- When an interactive job deletes the spooled file
- When the QSPLMAINT system job activates this process

The event is *not* triggered:

- When a Clear Output Queue (CLROUTQ) command is executed
- When a spooled file prints from the output queue
- When a batch job is used to CLROUTQ (unless the batch job also initiates a RCLSPLSTG *NONE command)
- When the subsystems start (unless the autostart job also initiates a RCLSPLSTG *NONE command)

For more information on these commands, refer to the *Work Management Guide*, SC41-5306.

10.11.3 Detecting Damage in QSYS Physical Files

If it is too difficult to schedule a reclaim storage process to detect damage to files, consider running the procedure provided in this section. This procedure detects for full or partial damage in physical files that are in the QSYS file system. Once compiled, the program does not look for damage in other file systems (for example, QDLS) or scan for damage to other object types.

The following example shows the program listing.

```
PGM (&START &END)
  DCL &START *CHAR 10      /* Library to start with */
  DCL &END *CHAR 10        /* Library to quit after */
  DCL &NBRRCD S *DEC 6
  DCLF      FILE(QSYS/QADBREF) ALWNULL(*YES)
  /* position in the file to the partial key lib name +
  (one key field) */
  OVRDBF    FILE(QADBREF) POSITION(*KEYAE 1 QADBREF +
  &START) OPNSCOPE(*JOB)
RCF:RCVF
  /* DONT FORGET TO CHECK FILES STARTING WITH $,@, # */
  /* START WITH $.. AND GO TO Z99999 */
  MONMSG   MSGID(CPF0864) EXEC(GOTO ENDP)
  /* block up to storage xfer (128k for risk, 32 k for cisc) */
  CHGVAR &NBRRCD S (32000 / &DBXRDL)
  OVRDBF FILE(&DBXFIL) OPNSCOPE(*JOB) +
  SEQONLY(*YES &NBRRCD S)
IF (&DBXLIB > &END) THEN(GOTO ENDP)
IF (&DBXATR=' PF'|&DBXATR=' TB') THEN(DO)
  IF (&DBXRDL > 4) THEN(DO)
    CPYF +
      FROMFILE(&DBXLIB/&DBXFIL) TOFILE(*PRINT) +
      FROMMBR(*ALL) FROMRCD(1) INCCCHAR(*RCD 1 +
      *EQ ' Q-*')
    MONMSG CPF0000
  ENDDO
```

```

ELSE DO
  /* a file with a length of 1 will still fail */
  CPYF +
  FROMFILE(&DBXLIB/&DBXFIL) TOFILE(*PRINT) +
  FROMMBR(*ALL) FROMRCD(1) INCCCHAR(*RCD 1 +
  *EQ ' -')
  /* File record length of 3 or less, will fail */
  MONMSG CPF0000
ENDDO
ENDDO
DLTOVR FILE(&DBXFIL)
GOTO RCF
ENDP:
ENDPGM

```

The job log shows the following error messages:

- Message CPC2957: No records copied from file &1 in &2.
This message indicates that the procedure does not detect problems with records in the file.
- Message CPD3244: File &1 in library &2 is damaged.
- No messages if the library does not have any files in it.

Note: The source listing is provided “as is,” on a best-effort basis for your use. You cannot use it in programs for resale. Compile and test it on your system before putting the procedure into production. For example, consider running this procedure in the interim of your scheduled RCLSTG processes.

Refer to information APAR I110690 for details on this process. For more information on accessing APARs, see *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01.

10.11.4 Detecting Damage in Physical Files

If it is too difficult to schedule a reclaim storage process to detect damage to files, consider running the Copy File (CPYF) command. The Copy File command performs an unblocked read of the input file. It notes errors if the command is entered in combination with the Override Database File (OVRDBF) command shown here:

```

OVRDBF FILE(library-name/file-name) SEQONLY(*YES 1)
CPYF FROMFILE(file-name) TOFILE(*PRINT) FROMRCD(1) ERRVL(999999999)

```

Using a value of “1” for the FROMRCD parameter causes a read by a relative record number. Specifying a 999999999 value for the errors allowed parameter ensures that records are not blocked. That is to say, if ERRVL is specified as “0” or *NOMAX, blocking may occur despite the SEQONLY parameter on the OVRDBF command. Specifying the parameters as shown in the example ensures that every record is read and that errors are logged if damage is encountered.

Run the CPYF process or the RCLSTG command after an unexpected failure such as a power or equipment outage or an abnormal termination of an application where files are not closed properly.

10.12 Power Down System

The Power Down System (PWRDWNSYS) command powers off and restarts (IPLs) the system. It is a straight forward function on systems prior to V4R1, where:

- Job logs are produced.
- The run priority does not change.
- The timeslice value does not change.

If your PWRDWNSYS is set to run at a scheduled time but does not, you need to:

1. Make sure the key lock switch on the system control panel is set to normal or auto.
2. Release the QSYSSCD job if it is held. If it is not running, stop then start cleanup. Or, if you do not want to use cleanup, specify "Y" to allow automatic clean up and specify *NONE for the time cleanup starts each day on the Change Cleanup Options (CHGCLNUP) display.

For V4R2 and later, be aware of the following considerations and options:

- How the Change Command Default (CHGCMDDFT) command affects commands such as PWRDWNSYS
- The timeout option parameter
- The end subsystem option parameter
- The time to terminate the system is shortened.

These considerations are described in the following sections.

10.12.1 Change Command Default Considerations

Since parameter levels override an element level on commands, you need to be aware of a change to the parameter structure for the PWRDWNSYS command. Beginning with V4R1, the PWRDWNSYS command uses a parameter list for the restart parameter. The restart type (RESTART) parameter is an element. Since CHGCMDDFT affects the parameter level, and *not* an element level, you must specifically code the restart type parameter value in your CL command if you choose values other than the defaults.

When using CHGCMDDFT to change the PWRDWNSYS command defaults, you can be misled when you see one of the following messages:

- The CPC6260 message: Default values changed for command PWRDWNSYS
- The CPD6265 message: List item changed, but SNGVAL at higher level

These messages are confusing, but true. Again, changing the parameter level overrides the change at an element level.

We recommend that you use the Change IPL Attributes (CHGIPLA) command if the IPL RESTART parameter is *SYS or *FULL. This way the default of *IPLA on the PWRDWNSYS command utilizes the desired parameter value for the next IPL.

10.12.2 End Subsystem Option

The end subsystem option (ENDSBSOPT) parameter controls the action taken when ending active subsystems. It has no effect on jobs already in the ending status. The ENDSBSOPT parameter can specify any of these actions:

- ***DFT**

The subsystems end with no special ending options, and job logs are produced. The run priority and timeslice value remain the same.

- ***NOJOBLOG**

Neither the subsystem jobs nor the user jobs in the subsystem produce job logs. This option can significantly reduce the time required for PWRDWN SYS to complete. However, if a problem occurs in a job, there is no job log to record the information, which makes problem analysis difficult or impossible.

Note: If OPTION(*IMMED) is specified, no job logs are produced during PWRDWN SYS regardless of the ENDSBSOPT value. However, these job logs are produced on the next IPL unless *NOJOBLOG is specified. The subsequent IPL may be faster, but the system does not power down more quickly if you enter:

```
PWRDWN SYS OPTION(*IMMED) ENDSBSOPT(*NOJOBLOG)
```

- ***CHGPTY**

The CPU priority of jobs that are ending changes to a higher value (that slows those jobs down). The remaining active jobs on the system, therefore, have better performance. However, jobs that are ending take longer to finish than otherwise. This option is ignored if the system is ending in a controlled fashion. If the *DELAY time limit expires, this option takes effect immediately.

- ***CHGSTL**

The timeslice of jobs that are ending changes to a lower value. The remaining active jobs on the system may have better performance when *CHGSTL is specified. However, jobs that are ending may take longer to finish. This option is ignored if the system ends in a controlled fashion. If the *DELAY time limit expires, this option takes effect immediately.

Note: The ENDSBSOPT is also available on the End Subsystem (ENDSBS) and End System (ENDSYS) commands.

The TIMOUTOPT and ENDSBSOPT parameters are shown in Figure 49 on page 163.

10.12.3 Timeout Options for Power Down System

The timeout option (TIMOUTOPT) parameter specifies the action to take when the system does not end within the time limit specified by the QPWRDWN LMT system value. If the system does not power down, it usually indicates that a job is hung, is in a loop, or did not end before the QPWRDWN LMT value was reached. If the time is exceeded, the next IPL is abnormal and, therefore, takes longer.

```

Power Down System (PWRDWN SYS)

Type choices, press Enter.

How to end . . . . . *CNTRLD *CNTRLD, *IM
Delay time, if *CNTRLD . . . . . 3600 Seconds, *NO
Restart options:
  Restart after power down . . . *NO *NO, *YES
  Restart type . . . . . *IPLA, *
IPL source . . . . . *PANEL *PANEL, A, B

Additional Parameters

End subsystem option . . . . . *DFT *DFT, *NOJOBLOG, *CHGPTY, *CHGCTL
+ for more values
Timeout option . . . . . *CONTINUE *CONTINUE, *MSD, *SYSREFCDE

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this disp
F24=More keys

```

Figure 49. The Power Down System Command

The actions specified are:

- ***CONTINUE**

The system ignores the timeout condition and continues powering the system down. If RESTART(*YES) is specified, the system restarts automatically. A minimum of information is available for service to debug the system.

- ***MSD**

The system issues a main store dump that is used by service to debug the system. It restarts after the dump is finished if the main store dump manager is configured correctly.

See *AS/400 Licensed Internal Code Diagnostic Aids—Volume 1, LY44-5900*, for more information on the main storage dump manager.

- ***SYSREFCDE**

The system displays a B9003F10 SRC and stops. This allows service to debug the system. On releases prior to V4R2, if the PWRDWN SYS OPTION(*IMMED) command takes longer than specified in the QPWRDWN LMT system value, a B9003F10 SRC code is displayed and the power down is stopped. For V4R2 systems, the SRC is “ignored” and the system continues to power down. To display the SRC and stop the power down, specify TIMOUTOPT(*SYSREFCDE) on the PWRDWN SYS command.

Note: Any changes to the default settings apply only to the next IPL.

Recommendation

If you have several Integrated PC Servers, the default value of 600 may not be enough for the QPWRDWN LMT value. Refer to Figure 8 on page 34 for a tool that you can use to show how long a PWRDWN SYS takes on your system.

10.12.4 Time to Terminate Improves System Availability

The PWRDWNSYS process and the ENDSYS and ENDSBS processes increase as more jobs use the system. Cleanup occurs when you terminate the system with these commands. Frequently the PWRDWNSYS, ENDSYS, and ENDSBS processes are followed by an immediate IPL, which repeats the cleanup process.

To reduce the unavailability caused by duplicate cleanup work, job termination:

- No longer de-allocates a job's associated communications device.
- Skips the clean up of temporary objects, since they are discarded on the subsequent IPL.
- Avoids closing sign-on files on a PWRDWNSYS.
- Avoids de-allocating storage pools.
- Cancels event monitors.

These changes improve the availability of the system.

10.13 Work Management APIs

For customers who want to further customize their operations, there are many APIs that are used for work management functions. Overhead is less when programs use APIs rather than the related (or in some cases, equivalent) CL command, so there is a performance advantage with using APIs as well.

This section discusses those work management APIs that are of latest importance, including:

- QUSRJOB Retrieve Job Information
- QUSLJOB List Job
- QWTCHGJB Change Job
- QUSCHGPA Change Pool Attributes

10.13.1 QUSRJOB Retrieve Job Information API

The QUSRJOB API retrieves specified information about a job. Fields available for use by programs incorporating this API include:

- The qualified job name
- Active job information
- Library list information
- Active job status
- Performance information
- WRKACTJOB information
- Job queue and output queue information
- Message logging information
- Job attribute information
- System pool identifier
- Number of auxiliary I/O requests (including database and non-database paging)
- Number of lock waits (internal, machine, database, and non-database) and the amount of time spent on the lock waits
- Response time total
- Purge information

QUSRJOB is enhanced to include a field that returns the system pool identifier in which the job is running. It is also updated to reflect thread information.

The system pool identifier is a key field beginning in V4R2 to identify the system pool in which the job is running. These identifiers are not the same as those specified in the subsystem description, but are the same as the system pool identifiers shown on the Work with System Status (WRKSYSSTS) display. The current system pool identifier returned by QUSRJOB1 is the actual pool in which the initial thread of the job starts.

The information accessed and produced by using this API allows for a customized analysis of work on the system. It can help control what happens when jobs do not start or perform as expected such as understanding or controlling resource contention (locks).

10.13.2 QUSLJOB List Job API

The QUSLJOB API generates a list of all or some jobs on the system, similar to the list produced by the Work with User Jobs (WRKUSRJOB) command. It is also updated to reflect thread information on V4R2 systems.

Fields available for use by programs incorporating this API include:

- The qualified job name
- The type of job
- Status of the job (for example, active, on the job queue, or on an output queue)
- System pool identifier

The system pool identifier is a key field beginning in V4R2 to identify the system pool in which the job is running. These identifiers are not the same as those specified in the subsystem description, but are the same as the system pool identifiers shown on the Work with System Status (WRKSYSSTS) display. The current system pool identifier returned by QUSLJOB is the actual pool in which the initial thread of the job runs.

The information accessed and produced from this API allows for a customized analysis of work on the system. You can use it to help in resource or performance analysis.

10.13.3 QWTCHGJB Change Job API

The Change Job API changes a list of attributes about a job similar to the Change Job (CHGJOB) command. The current value of most job attributes can be retrieved with other APIs such as:

- QUSLJOB—List Job
- QGYOLJOB—Open List of Jobs
- QWCRTVCA—Retrieve Current Attributes
- QUSRJOB1—Retrieve Job Information

Fields available for use by programs incorporating the QWTCHGJB API include:

- Qualified job name
- Job change information
- Device recovery action
- Job accounting information
- Job message queue full action
- Printer device name
- Schedule date and time
- Library list
- Output queue name

- Logging level information
- Run priority
- Maximum time to wait for an instruction to complete
- Job switches

The information accessed and produced by using this API can be used to help manage jobs more effectively.

10.13.4 QUSCHGPA Change Pool Attributes API

The QUSCHGPA API allows a change of tuning parameters for system pools similar to those that are changed by the Change Shared Storage Pool (CHGSHRPOOL) command. Depending on whether the base pool, shared pool, or private subsystem pool is to be changed, QUSCHGPA issues the appropriate command.

The tuning parameters that you can affect are:

- Storage pool size
- Activity level
- Paging options
- Minimum and maximum page faults
- Page faults per thread
- Pool priority
- Minimum and maximum pool size

The information is similar to the information that is available on the Work with Shared Storage Pools (WRKSHRPOOL), Work with System Status (WRKSYSSTS), and Work with Subsystems (WRKSBS) displays. Use the QUSCHGPA API to tune storage pools and paging options interactively without having to know which subsystem monitor allocated the pool. You can also develop a performance adjustment process much the same as the one provided with the QPFRADJ system value. Refer to Section 8.3, “What a User Can Do to Influence System Performance” on page 100, for more information on performance adjustment.

In some cases, pool changes do not take effect immediately. A save or restore operation might be using some of the storage allocated to a pool, for example, or the system might be using some of the storage allocated to the base pool (pool number 2). The size is changed only when the storage being used is free again.

Refer to the *System API Reference*, SC41-5801-01, for more information on these and other APIs.

Chapter 11. Availability and the PTF Process

One cause of unavailability for AS/400 systems is program temporary fixes (PTFs). PTFs “cause” unavailability from two aspects:

1. Downtime has to be scheduled to apply fixes on a preventive basis
2. Downtime is at times caused by not having fixes applied to prevent problems (assuming a fix is available).

IBM has made improvements in how PTFs are applied to decrease the downtime involved. This chapter addresses how the process has been improved and also provides useful information for developing a backup and recovery strategy. This chapter is helpful to the system administrator for developing a change and system management plan.

Topics addressed in this chapter include:

- PTF terminology
- Preventive service planning
- Applying and removing PTFs and IPLs
- Product level support
- Conditional PTFs
- Cumulative PTF package
- Applying PTFs for the next release prior to installing it
- Distributing PTFs
- Requesting PTFs and cover letters using the Internet

11.1 PTF Terminology

We begin the discussion with a brief introduction to the terminology used in connection with PTFs.

- PTF

A PTF is a set of program modules meant to replace an earlier version of a program module. PTFs are first loaded onto the AS/400 system as save files.

- Superseded PTF

A superseded PTF has its program modules replaced by a PTF that was released at a later time. The superseded PTF *is not* installed on the system. The superseded PTF *is* included in a PTF that was produced later and *is* installed on the system.

- Prerequisite and Co-requisite PTFs

Prerequisite PTFs must be applied before or at the same time as the PTF, which requires them, is applied. Co-requisite PTFs must be applied at the same time as the named co-requisite PTF that needs them.

- PTF Cover Letter

Each PTF consists of program modules and a text file that describes both the problem fix and any activation instructions to affect the installation of the PTF. The cover letter also documents co-requisite PTFs, prerequisite PTFs, superseded PTFs, and PTFs that are distributed together.

- RETAIN

RETAIN is the worldwide database for IBM problem management and the central repository of PTF files. You access this database with any and all PTF orders, no matter which ordering method is used.

- Preventive Service Planning (PSP)

Preventive service planning (PSP) is a collection of information regarding PTFs applicable to AS/400 hardware and software problems. PSP information should be investigated on a regular basis by using ECS or the Internet. See Section 11.2, “Preventive Service Planning” on page 169, for information about receiving PSP buckets.

- High Impact or Pervasive (HIPER) PTF

High impact or pervasive (HIPER) PTFs correct serious system problems. These problems usually relate to data integrity, system, and communication processing problems that the IBM AS/400 engineers suspect all or almost all customers experience. These PTFs should be proactively applied. You can find them by reviewing the PSP package or using the ALERT/400 service offering.

- Cumulative (CUM) PTF package

Cumulative (CUM) PTF packages are a regularly released set of PTFs that contain all the HIPER PTFs and selected non-HIPER PTFs. The selection criteria changes over time, but generally the PTFs that have had a certain number of downloads through ECS are selected for the next CUM package. Refer to Section 11.6, “Cumulative PTF Package” on page 176, for more information on CUM packages.

- Electronic Customer Support (ECS)

Electronic customer support (ECS) is an integrated set of functions designed to help service and support the AS/400 system. ECS provides:

- Hardware and software problem analysis, reporting, and management
- Copy screen image
- Question and answer support
- IBM technical and product information access

See *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01, for detailed information on problem reporting using ECS.

- Immediate versus Delayed PTF

Some program modules cannot be replaced while in use. Immediate PTFs are those that can be activated while the system is up and performing normal work, where delayed PTFs are activated during a normal IPL. Delayed PTFs can be loaded in groups during a single apply function. Some immediate PTFs do not require an IPL but do require that the product they affect is not in use. Some require a restart of the IOP or a vary off and on of a network description.

- A-side and B-side of Licensed Internal Code (LIC)

On the AS/400 system there are two copies of Licensed Internal Code (LIC) known as the A-side and B-side. When a LIC PTF is temporarily installed, it is only installed on the B-side. When a LIC PTF is applied permanently, it is applied to both the A-side and B-side copies of the LIC. If all the LIC PTFs are applied permanently, the system can only be IPLed on the A-side.

Note: Most AS/400 systems should use the B-side copy of the LIC unless otherwise directed by IBM support representatives.

- Temporary versus Permanent PTF

When a PTF is temporarily applied, the original program module is renamed and the new module replaces the original one. If the new module proves to be defective, the original module can be reactivated simply by removing the new PTF containing the program module with the Remove PTF (RMVPTF) command.

When a PTF is permanently applied, the original module is destroyed. There is no easy way to go back and use the old module. Nonetheless, there are two reasons for applying a PTF permanently:

- To save disk space
- A new PTF requires the older one to be permanently applied

11.2 Preventive Service Planning

We recommend that you proactively review what fixes are available to resolve problems for any of the products and options installed on your AS/400 system. The Preventive Service Planning (PSP) package can be used for this purpose. PSP information is available for each cumulative PTF package, Licensed Program Products, Vertical Licensed Internal Code (VLIC), and Horizontal Licensed Internal Code (HLIC). Contracted users can access Preventive Service Planning (PSP) information by:

1. Using SNDPTFORD to send a pre-determined PTF order number and reviewing the returned member in the QGPL/QAPZCOVER file, or
2. Locating the following URL on the Internet:

<http://www.AS400Service.ibm.com/as4sde/sline003.NSF/>

Cover letters that have been downloaded can be reviewed by entering:

```
CPYF FROMFILE(QGPL/QAPZCOVER) TOFILE(QGPL/QPRINT) FRMMBR(QSF98vrm)
```

In this command, vrm is the Version/Release/Modification number for which you requested the PSP information (for example, 420 for V4R2, as listed below).

The predetermined PTF order numbers containing this PSP information and summary list of HIPER PTFs and fixes by release are:

<u>PSP Order Number</u>	<u>Version and Release</u>
SF98420 MF98420	Version 4 Release 2.0
SF98410 MF98410	Version 4 Release 1.0
SF98370 MF98370	Version 3 Release 7.0
SF98360 MF98360	Version 3 Release 6.0
SF98320 MF98320	Version 3 Release 2.0
SF98310 MF98310	Version 3 Release 1.0

11.3 Applying and Removing PTFs and IPLs

To progress toward 24 x 365 system availability, it is necessary to reduce or eliminate maintenance operations that cause the system to be off-line. These days, a primary cause of downtime is the application of PTFs, especially fixes that are applied to the Licensed Internal Code (LIC). More and more PTFs need to be, and are, built to apply immediately to help reduce downtime.

When applying a microcode PTF on systems prior to V4R1, the system performs an IPL from the A-side to the apply PTF stage. That includes up to the C6nn nnnn set of System Reference Codes (SRCs). Microcode PTFs are applied. Before OS/400 is started, the system powers down and performs IPL again using the B-side. The IPL is selected using the control panel or the defaults of the PWRDWSYS command. This process is fairly lengthy, and requires the LIC to be loaded for each IPL.

On V4R1 and later systems, if the microcode PTFs do not affect the MFIOP microcode, the system does not reload the MFIOP when applying the PTFs. This process avoids the need for a second IPL.

For example, a user may apply a delayed microcode PTF while on the B-side, and issue a PWRDWSYS with a RESTART(*YES) and TYPE(*SYS). The system powers down and performs a mini-IPL without reloading the code into the MFIOP. The microcode PTF subsequently is applied during the SLIC part of the IPL. After the PTF is applied, the system IPLs again without reloading the MFIOP code. If the PTF affects the MFIOP code, the MFIOP code is reloaded during the IPL.

The advantage of this approach is that the system does not have to IPL all the way up to OS/400 on the A-side, before coming down and performing IPL again on the B-side. This way, the IPL time is considerably reduced, and the PTF application is faster than it was prior to V4R1. This process is pictured in Figure 50.

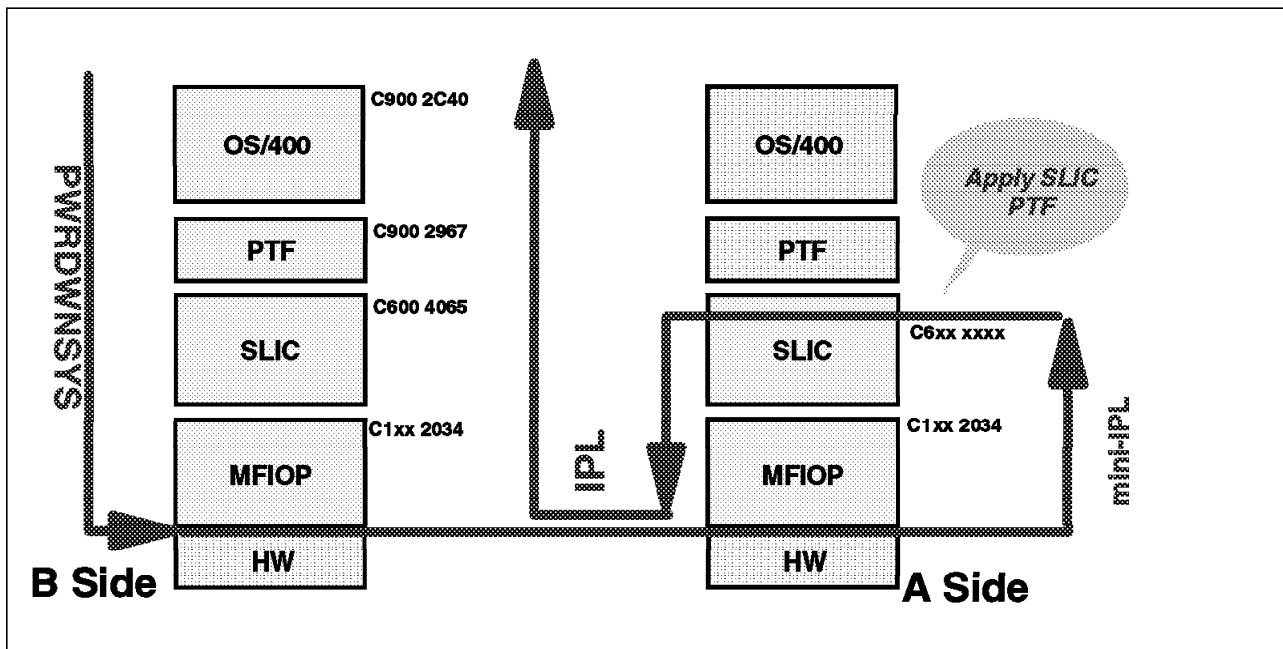


Figure 50. A View of the Apply PTF Process

11.3.1 LIC PTF Apply and System Availability

To apply LIC PTFs, the system must be able to access the A-side. If the system IPLs from the B-side, a message appears indicating the wrong copy of LIC is in use. To apply LIC PTFs, we recommend that you use the following command:

```
PWRDWSYS OPTION(*IMMED) RESTART(*YES) IPLSRC(B)
```

This command performs an IPL to the A side to apply the LIC PTFs and then to the B side to apply the rest of the PTFs. Since the Job Scheduler can only perform IPLs on the B-side, LIC PTFs are not applied. If these LIC PTFs are the prerequisite of any OS/400 licensed program PTFs, the latter is not applied.

To increase the number of immediate PTFs for LIC, the apply procedure for nucleus modules was changed to fixes for nucleus modules beginning with V4R2. Space usage allows more PTFs to be built for an immediate apply. Changing the apply process for nucleus modules and minimizing the space usage enable the building of more PTFs as immediate, thus shortening or completely removing downtime caused by IPLs required to apply fixes.

Note: All the changes presented in this section are for RISC technology only. They do not apply to any CISC releases.

11.3.2 Applying PTFs without an IPL

Beginning in V3R1, PTFs are built to produce PTFs without requiring IPLs to apply them whenever possible. PTFs that can be applied immediately have additional activation steps, which are described in the cover letter. Updateable and non-updateable actions exist. Updateable action PTFs are shipped with an action-exit program to verify the actions required to activate the PTF. The action-exit program is called by the DSPPTF command to list the actions necessary and verify whether the actions are completed. The exit program is called when the details for the PTF are shown using option 5 on the DSPPTF status panel or when you specify SELECT(*ACTRQD) on the DSPPTF command. The status of the PTF is updated if the actions are complete. The status changes from *Temporarily applied—PND* to *Temporarily applied* after the required actions are complete.

Non-updateable action required PTFs do not have exit programs to verify that the actions have been done. These PTFs remain in a *Temporarily applied—ACN* status until the next IPL when the status changes to *Temporarily applied*.

11.4 Product Level Support

A product level PTF is a logical way of grouping together all fixes for a specific solution. Product levels are used to simplify the prerequisite and co-requisite processing. They can be regarded as an extension to the Version/Release/Modification convention.

Product levels can be applied and removed as a group (within a product). One PTF number is assigned for a particular solution group for each release. This is not the same as a cumulative PTF package. The product level PTF is updated more often, and contains PTFs only for certain products which are interrelated. It does not contain fixes for all of OS/400.

The product level PTF concept is implemented within IBM's RETAIN system when an order is received and the fixes are packaged. PTF management support for AS/400 systems does not address product level PTFs. The system administrator

should be aware of levels when ordering and installing fixes. There is no need for an end user to know what a product level is.

The PTFs that form the product-level PTF are shipped to you by mail on a CD-ROM or a tape. The files are not available to download directly through Electronic Customer Support (ECS) because the files are too large and take a long time to load. The *only* way to receive a product-level fix is on media. To order, use the Send PTF Order (SNDPTFORD) command on your AS/400 system if you use electronic customer support (ECS) or contact your local service provider to request a media shipment.

Product-level PTFs can also be thought of as a variation on or an extension to an additional function PTF that provides optical support at V3R6 with feature code 1980. Product-level PTFs can add a new function to the operating system or microcode without upgrading either. Product levels offer the following features:

- Allow delivery of new function with minimal impact to the customer
- Allow delivery of new function without the overhead of installing a new release
- Allow delivery of new function without the overhead of producing a new release (a direct benefit to IBM and an indirect benefit to the customer)
- Remove unnecessary workload and DASD requirements for customers not wanting new functions
- Provide a single PTF for both base and product levels of code
- Eliminate save and restore compatibility problems

A product level consists of previous level parts and PTFs for those parts. A product level is designated with a two-position alphanumeric field, starting from a value of 00 and incrementing in steps of 10. A cumulative PTF package contains PTFs for all product levels. The product level concept only applies to OS/400, LIC, and the PTFs for these products—not licensed programs or operating system options. By default, the system installs the first level found that is equal to or greater than the currently installed level.

To find out the product level of an AS/400 system, enter the DSPPTF command or select option 10 on the LICPGM menu to *Display installed licensed programs*. The following displays show examples of how the product level is appears.

```

Display PTF Status
System: SYSTEMXX
Product ID . . . . . : 5769999
IPL source . . . . . : ##MACH#B
Release of base option . . . . . : V4R2M0 L00
Type options, press Enter.
  5=Display PTF details  6=Print cover letter  8=Display cover letter

PTF
Opt ID Status IPL
Action
- TL90004 Temporarily applied None
- TL90003 Superseded None
- TL90002 Permanently applied None
- TL90001 Permanently applied None
  5 MF17331 Temporarily applied None
- MF17291 Temporarily applied None
- MF17272 Temporarily applied None
- MF17271 Superseded None
- MF17244 Temporarily applied None

More...
F3=Exit F11=Display alternate view F17=Position to F12=Cancel

```

Figure 51. How to select the Display to See the Product Level

To determine the lowest and highest level of the product that this PTF can install on, look at the minimum and maximum system levels on the second page of the General Information display relating to the PTF, as shown in Figure 52.

```

General Information
Product ID/PTF ID . . . . . : 5769999 MF17331
Release . . . . . : V4R2M0

Target OS/400 Release . . . . . : V4R2M0
Minimum-Maximum level . . . . . : L00-L00

Bottom

Press Enter to continue
F3=Exit F12=Cancel

```

Figure 52. The Minimum and Maximum Levels Allowed for a Single PTF

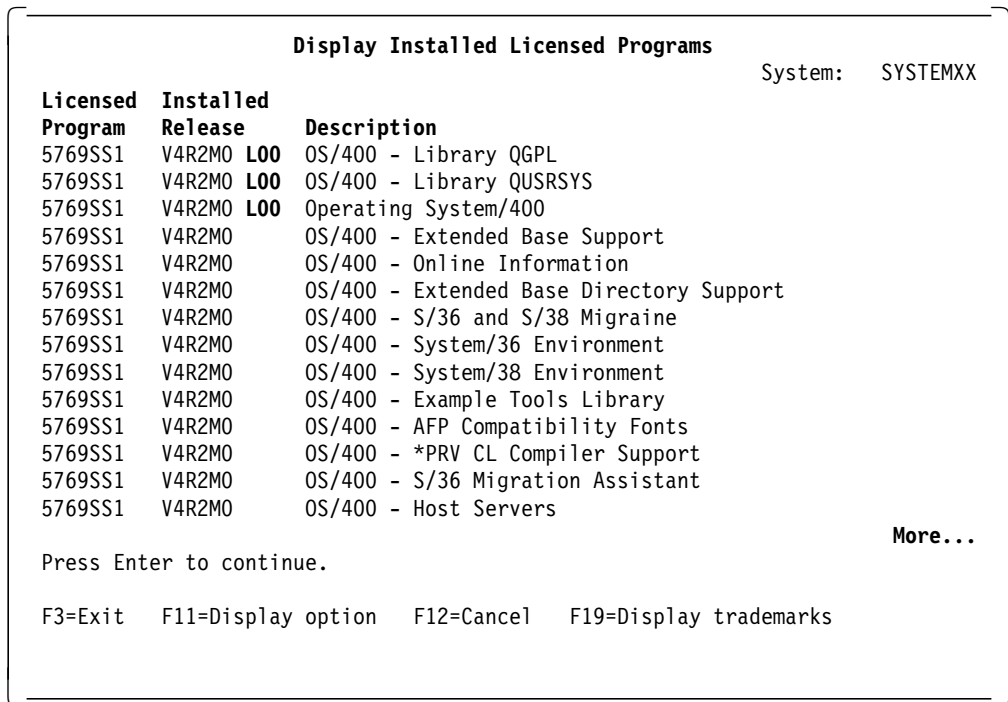


Figure 53. Levels Shown on a Display Installed Licensed Programs Display

The following tables list some of the product-level PTFs available at the time of writing this redbook:

Table 12. Windows Product Level PTFs

OS/400 Release	Windows 95 Client	Enhanced Windows 3.10 Client	Windows 3.1 Client R311
V4R2	not available	SF99053	SF99054
V4R1	SF99049	SF99052	SF99048
V3R7	SF99043	SF99051	SF99047
V3R6	SF99042	not available	SF99046
V3R2	SF99041	SF99050	SF99045
V3R1	SF99040	not available	SF99044

Table 13 (Page 1 of 2). Product Level PTFs for Miscellaneous Functions

OS/400 Release	Product	Group PTF
V4R2	Database	SF99102
V4R1	Database	SF99101
V3R7	Database	SF99100
V4R1	Internet Connection Server	SF99086
V3R7	Network Station	SF99081
V3R2	Network Station	SF99080
V4R2	Backup Recovery	SF99074
V4R1	Backup Recovery	SF99073
V3R7	Backup Recovery	SF99071

<i>Table 13 (Page 2 of 2). Product Level PTFs for Miscellaneous Functions</i>		
OS/400 Release	Product	Group PTF
V3R6	Backup Recovery	SF99070
V3R2	Backup Recovery	SF99072
V3R6	Optical Support	SF99087

You can only order these product level PTFs by using ECS. Specify SNDPTFORD PTFID(SF99nnn) DELIVERY(*ANY) when ordering.

11.4.1 DB2/400 PTF Strategy

To help clarify the PTF strategy for DB2/400, this section discusses product level PTFs and Fixpacks.

PTFs for DB2/400 are ordered as a *Fixpack*. A *DB2/400 Fixpack* is a method of bundling DB2/400 PTFs together into a single package. Also known as product level PTFs, the *Fixpack* automatically includes all required prerequisite and co-requisite PTFs, packaged on a single media similar to traditional CUM packages. This packaging allows you to order the most recent, fully tested DB2/400 PTFs with one request.

A new Fixpack is released every eight weeks. Outside this packaging method of delivery, you can order PTFs individually.

11.5 Conditional PTFs

Conditional PTFs are known as either co-requisite or prerequisite PTFs. Prerequisite checking occurs when the PTF with the prerequisite is set for apply. The PTF with the prerequisite knows about its prerequisites, but the prerequisite PTF itself does not know what PTFs depend upon it. When a cumulative PTF package is set for apply, use the GO PTF command and select option 8. Any LIC PTFs that are considered prerequisites for PTFs on the cumulative tape are set for permanent apply when the dependent PTF is set for apply.

Note: The action of setting the dependent PTF for apply sets the LIC prerequisite PTF for a permanent apply.

On V4R2 systems, co-requisite PTFs are applied and removed as a group. Prerequisite and co-requisite requirements are enforced.

On systems prior to V4R2, it is possible that only some of the PTFs required to fix a problem are applied. Prerequisite checking is done when the PTF with the prerequisite requirement is set for an apply. The system administrator must closely review the PTF cover letter, job log of the apply process, and status of all applied PTFs to ensure all prerequisite and co-requisite requirements are met to ensure a satisfactory installation of the PTF.

Co-requisite information is contained in the PTF module itself and verification is done by the PTF apply job. If a required PTF is not available, the PTF to be applied is not processed. On systems prior to V4R2, this information is only available in the cover letter. Cross checking the information to make sure all PTFs are loaded is the responsibility of the user.

If the PTF is for an MFIOP module, the user has to perform several IPLs before all the prerequisite and co-requisite PTFs are loaded and successfully applied.

On systems prior to V4R2, follow these steps:

1. Order PTF A.
2. Load and apply PTF A (with DELAYED(*YES)).
3. Perform an IPL.
4. Check the job log and the PTF status. Determine if PTF A was not applied because one or more co-requisite PTFs are missing.
5. Go back to step 1 to order the appropriate co-requisites.

On V4R2 systems, follow these steps:

1. Order PTF A.
2. Load and apply PTF A (with DELAYED(*YES)).
3. Receive an error that co-requisite PTF B is missing.
4. Order PTF B.
5. Load and apply the PTF B (with DELAYED(*YES)).
6. Perform an IPL.

Conditional PTF prerequisites and co-requisites are enabled at a product level. That is, PTF A can have different PTFs as prerequisites or co-requisites depending on the product level of PTF A.

Note: Co-requisite PTFs are sometimes referred to as group PTFs. However, the “group” label is misleading as it implies that the fixes are somehow packaged together and can be acted on as a group. This is not true for systems prior to V4R2.

11.6 Cumulative PTF Package

Not all PTFs produced are included in a cumulative PTF package. The following criteria help determine if a given PTF is incorporated into a subsequent cumulative package:

- The number of customer orders for a particular PTF
- The severity of the associated APAR
- Whether the PTF is designed as HIPER
- Problems that crash a system and require an IPL
- High availability options (mirroring or checksum, for example)
- Problems that require a reinstall or restore
- Save and restore problems that prevent data recovery
- Problems affecting installability of a product or feature
- Problems affecting a product’s major useability function
- Problems requiring an unnecessary replacement of hardware parts
- Pervasive customer satisfaction problems
- Data integrity and security problems
- If the PTF has a corresponding PTF in a previous release package (to minimize the rediscovery)

PTFs that have special instructions, which might break the system if they are not followed, are not usually incorporated into the cumulative package.

11.6.1 Applying CUM Packages

Cumulative packages can be applied at the same time as individual PTFs.

If you receive the individual PTFs on tape, follow these steps:

1. Place the tape containing the individual PTFs in the tape drive.
2. Enter GO PTF and select option 8 (Install program temporary fix package).
3. Specify Automatic IPL N on the Install Options for Program Temporary Fixes display.

If you receive the individual PTFs through ECS, perform these steps:

1. Enter GO PTF and select option 8 (Install program temporary fix package).
2. Specify *SERVICE for the device parameter, and select N for Automatic IPL.

When the GO PTF command is used to install PTFs from *SERVICE, all PTFs in *SERVICE not already installed are installed.

*SERVICE means that the PTF exists in a save file in the QGPL library and that the system checks for these PTFs during the installation process.

11.7 Applying PTFs for the Next Release Prior to Installing the Next Release

PTFs for a higher release can be downloaded to a previous release system and applied if the receiving system is at R310 or later. The PTFs can be loaded and applied to the intended system after the system is upgraded to the release that the downloaded PTFs are for. Follow these steps:

1. Install the higher release (for example, R420).
2. Use the Load PTF (LODPTF) command to load the PTFs.
3. Enter *SAVF as the device name and QSFnnnnn for the SAVF name, or QMFnnnnn in the QGPL library.

For example, if PTF MF12345 for R420 is downloaded while you are still at R370, use the following command to load PTF MF12345 *after* installing R420:

```
LODPTF LICPGM(5769999) DEV(*SAVF) SELECT(MF12345)
SAVF(QGPL/QMF12345)
```

You can use the Apply PTF (APYPTF) command and apply the PTFs.

Note: The APYPTF command must be used instead of the GO PTF command with *SERVICE because the downloaded PTFs do not have an entry in the PTF index.

11.8 Distributing PTFs

Loading PTFs distributed from other systems can be done as a group. To do so, make sure the following procedure using the Send Network File (SNDNETF) command is implemented for distributing the PTFs:

1. Use SNDNETF to send the PTF save files to the remote system.
2. On the remote system, use the QPZGENNM API to generate a valid save file name for the PTF save file.
3. Create a save file in QGPL by that name and receive the PTF into that save file.
4. Use the QPZLOGFX API to log the PTF into *SERVICE.

Now the PTF is recognized by the system and is available for any PTF functions, such as LODPTF DEV(*SERVICE).

Refer to the *System Programmers Interface Reference*, SC41-5801-01, for information on how to use APIs.

Note: As an alternative method, use the SNDPTF command, which is available with the System Manager/400 licensed program product. This SNDPTF command sends the PTFS, loads them, and sets them for an apply with a single command on the remote system.

Refer to *System Manager Use*, SC41-5321, and *Managed System Services for AS/400 Use*, SC41-3323, as well as Section 9.8, "IBM SystemView System Manager for AS/400" on page 128 for more information on IBM SystemView System Manager for AS/400.

11.9 Requesting PTFs and Cover Letters Using the Internet

Internet access is an alternative connection method between the AS/400 system and the RETAIN system that houses PTF information. On the Internet, users can:

- Browse PTF information
- Request shipment of a CUM package
- Request shipment of selected PTFs

Enabling Internet access over using ECS to request delivery has an advantage for many customers in the U.S. due to the cost, availability of modems, and higher speed transfer. Internet access provides a graphical interface and has the potential for faster access and download of PTF information. This is a result of higher speed modems that are available through some Internet service providers and the possible use of T1 lines. ECS PTF delivery is limited to a SDLC 19.2 kbps synchronous connection.

For some customers, ECS remains the preferred method to request PTFs. In other cases where ECS is not available, Internet access provides a faster and more efficient solution than delivery of the PTF information through the "mail."

You can read PTF cover letters on the Internet. Cover letters are downloadable from the URL:

<http://www.as400service.ibm.com/as4sde/as4ptf.nsf/ptfbyrel>

Internet delivery does not replace or affect the traditional PTF order process. If you do not have access to the Internet, continue using the ECS PTF order function with the SNDPTFORD command.

To print cover letters for PTFs for a product installed on your system, enter the command:

```
DSPPTF LICPGM(57nnxxx) SELECT(yyyyyy) COVERONLY(*YES) OUTPUT(*PRINT)
```

In the command, xxx is the licensed program number, and yyyyyy is the PTF ID. The output can be viewed in the spooled file named QSYSPRT.

To print a cover letter of a PTF for a product not installed on your system, enter the following command:

```
CPYF FROMFILE(QGPL/QAPZCOVER) TOFILE(QGPL/QPRINT) FROMMBR(Qxxxxxxxxy)
```

In this command, xxxxxx is the PTF ID, and yy is the language ID. English PTFs do not have a language ID.

11.9.1 Media or PTF Cover Letter Order Scenario

This section describes the major events that occur for an Internet request to ship a CUM tape or PTF. The main steps in the process, the related program elements on the PC, and the AS/400 system, and their relationship are included. Follow these steps:

1. The request for a PTF order is initiated from the workstation by the user using a Web browser. The user is required to access the IBM service home page before an order can be initiated. The address of the IBM service home page for AS/400 Preventive Service Planning is:

<http://www.AS400Service.ibm.com/as4sde/sline003.nsf/sline003home>
2. To access PTF information, use the HTML pages to obtain all PTF cover letters and PSP buckets.
3. Initiate an order transaction by following the applicable instructions on the HTML pages.
4. To authenticate a user, you are required to provide a valid user ID and password.

11.10 For More Information

For more information on managing PTFs, refer to *Basic System Operation, Administration, and Problem Handling*, SC41-5206; *System Operation*, SC41-4203; and the following URL address:

<http://www.as400service.ibm.com/as400/service.html>

Chapter 12. Communications Error Recovery and Availability

A critical aspect of availability is the stability of the network used to connect with the various platforms in a business. The affect of error recovery procedures often overlaps other areas of the system. Therefore, it is essential that these communications errors cause as little impact to each platform as possible. However, due to the cost of redundant solutions within a network and to some network outages beyond our control, not everyone can have a fully fault tolerant network. This chapter discusses changes made to the AS/400 system to reduce the impact of communications error recovery to users on the system.

Communications errors should not require the system to be IPLed for recovery. The most severe action should be to vary off the affected configuration objects, then vary them on again. Error recovery should have minimum to no impact on unaffected users on the system.

Topics discussed in this chapter include:

- What communications error recovery procedures are
- Improvements on V4R2 systems
- Improvements on V4R1 systems
- Configuration tips and techniques
- Testing error recovery

12.1 What Communications Error Recovery Procedures Are

Communications error recovery procedures (ERP) are the set of actions that occur within a network or system when the connectivity between the host (AS/400 system) and any of its communication interfaces are interrupted. For example, communications ERP may apply to a router failure, pulled cable, power glitch, software malfunction, or any failure that can cause the affected communications session to enter into error recovery procedures. Error recovery is complete only when all users are successfully back online doing productive work.

The object of communications ERP is to contain the errors and limit the affect that a single point of failure has on other processes in the system.

The scope of communications error recovery involves far more than just reporting the error. It includes all work performed on the system for users to connect and get their working environments back up and operational. This includes the error reporting path, the de-activation path, and the activation path within the communications structure. In addition, other system work (starting and ending system jobs, for example) plays a significant role in communications error recovery processing.

12.2 Improvements on V4R2 Systems

With improvements in V4R2, the scalability of the impact on multi-processor systems allows minimal negative impact to other work in the system. Scalability occurs with the reduction of errors and the elimination of throughput bottlenecks for more parallel processing as the number of users increases.

Changes include:

- QCMNARB system value settings
- The addition of high performance routing in an APPN environment
- Job-end improvements
- Improvements in logging and service functions

Except for job ends, these topics are further described in the sections that follow. Ending jobs is discussed in 10.9, “End Job Abnormal” on page 152.

12.2.1 QCMNARB System Value Setting

Work that was done in the QSYSARB system job for APPC is moved out of QSYSARB and into communications arbiter jobs. These communications arbiters start when the system is IPLed. They are named QCMNARB01 through QCMNARB12.

There is one QCMNARBnn job per processor on the system. For example, on a 12-way processor, there are jobs named QCMNARB01 through QCMNARB12. The exception is that single processor systems have QCMNARB01 and QCMNARB02.

In addition, the APPN automatic creation, deletion, and vary on and off processing of APPN controllers and devices that was done by the QCLUS system job is now performed in the communications arbiter jobs.

Note: Throughout this chapter, the reference to APPN controllers and devices indicates objects created when the APPC controllers and devices indicate APPN(*YES).

The QCMNARB system value specifies the number of communications arbiter jobs that are available to process APPC communications. The QCMNARB setting can affect the performance at startup, takedown, and during error recovery for APPC communications and APPN autoconfiguration. The first time an APPC controller is varied on after an IPL, it is assigned to a communication arbiter job. On systems with many APPC controllers and devices, more QCMNARB jobs can result in improved APPC communication performance.

If QCMNARB is set to zero, the system functions as systems prior to V4R2. That is, the work is performed in QSYSARB and QCLUS, not in the communication arbiters.

Recommendation

Set the QCMNARB system value to a non-zero setting. It is also acceptable to leave it at the default value of *CALC, which typically results in a non-zero setting.

On multi-processor systems, a setting of *CALC results in one communication arbiter job for each processor on the system.

The system maintains a count of objects for each QCMNARB job and attempts to assign controllers to the communication arbiter job that is serving the *least* number of objects. This balances the amount of work that each QCMNARBnn arbiter job must perform. All devices that are attached to an APPC controller are serviced by the arbiter job to which the controller is assigned.

To determine which communication arbiter is assigned to a particular APPC controller, use the Display Controller Description (DSPCTLD) command and locate the QCMNARBnn system job that handles recovery for that controller and its devices. An example is shown in Figure 54 on page 183.

```

                                Display Controller Description
                                SYSTEMXX
                                02/27/98 20:27:22
Controller description . . . . . : CTLD      CTLD0001
Option . . . . . : OPTION    *BASIC
Category of controller . . . . . :          *APPC

Link type . . . . . : LINKTYPE *LAN
Online at IPL . . . . . : ONLINE  *YES
Character code . . . . . : CODE    *EBCDIC
Maximum frame size . . . . . : MAXFRAME 16393
Current maximum frame size . . . . . : 521
Remote network identifier . . . . . : RMTNETID SYSTEMSXP
Remote control point . . . . . : RMTCPNAME SYSTEMDMT
Initial connection . . . . . : INLCNN  *DIAL
Dial initiation . . . . . : DIALINIT *LINKTYPE
Switched disconnect . . . . . : SWTDCS  *NO
Data link role . . . . . : ROLE     *NEG
LAN remote adapter address . . . . . : ADPTADR 08005A202722
LAN DSAP . . . . . : DSAP        04
LAN SSAP . . . . . : SSAP        04
Autocreate device . . . . . : AUTOCRTDEV *ALL
System job . . . . . : QCMNARB01
Text . . . . . : TEXT      AUTOMATICALLY CREATED BY SYSTEM

Press Enter to continue.
                                Bottom

F3=Exit  F11=Nondisplay keywords  F12=Cancel

```

Figure 54. Communication Arbiter Job Assigned to the Controller

If there are APPC controllers that have non-APPC devices attached to them, the recovery for that controller and its devices continue to execute in QSYSARB. The same applies to V4R2 systems. That is, only APPC controllers with APPC devices attached take advantage of the multiple communication arbiters function.

12.2.2 High-Performance Routing

On V4R2 systems, APPN High-Performance Routing (HPR) functions are available. High-Performance Routing can provide many benefits for error recovery. One benefit is the ability to route around failed links. HPR can also avoid session failures by setting timers to ride out short-term outages.

HPR no longer combines half sessions at each intermediate routing node to provide a session from the source to the destination. To address enabling high-speed lines with fewer errors, HPR removes the requirement to implement protocols on a hop-by-hop basis. Instead, HPR implements all protocols on an end-to-end basis. This improves performance by reducing the number of protocol flows on each link.

Automatic network routing significantly improves performance in intermediate nodes. Functions, such as link-level error recovery, segmentation, flow control, and congestion control, are no longer performed in the intermediate nodes. Instead they are performed at the end points. This saves on storage on the intermediate node systems.

The following network attributes were added when High-Performance Routing was added. However, the system value settings can affect performance during communications error recovery for non-HPR, APPN configurations, and an HPR environment.

The HPR related network attribute values are:

- ALWVRTAPPN

The allow virtual APPN support network attribute is used to choose whether non-HPR APPN devices should be attached to the real APPN controller or to a virtual controller.

The greatest benefit of ALWVRTAPPN is that it can *eliminate* the need to automatically create and delete APPN device descriptions. When multiple routes are designed in an APPN network, there are multiple device descriptions—one for each possible route in the network. ALWVRTAPPN eliminates the overhead of the large amount of configuration objects, which results from the possibility of multiple routes.

Another reason you may want to change the ALWVRTAPPN value to *YES (the default is *NO) is to limit the number of devices that go through error recovery when a failure occurs. When virtual APPN controllers are used, only the real APPN controller goes into error recovery. Since the devices are attached to the virtual APPN controller, they do not go through the traditional error recovery work flow.

- VRTAUTODEV

When ALWVRTAPPN(*YES) is specified in the system network attributes, the virtual controller autcreate device (VRTAUTODEV) network attribute parameter indicates the maximum number of devices automatically created for each virtual controller. The VRTAUTODEV parameter can be used to control the upper limit for the number of autcreated APPN devices on APPN virtual controllers. The default value for this parameter is 100. This means for every 100 new APPN locations with which your system communicates, a new virtual APPN controller is automatically created.

Use the CHGNETA command to alter these settings.

Note: Each varied on virtual controller description is managed by a task on the system. Having multiple virtual controllers descriptions can allow for some parallel processing to occur. It also cuts down on certain queue lengths, such as the queue used to maintain the list of devices that are managed by the APPN virtual controller task.

Think about how much work is placed under one controller. Too much work under one controller can result in slower performance during connect, disconnect, and error recovery processes. Use a smaller number if your sessions last a short time.

A view of the WRKCFGSTS display with and without HPR is shown in Table 14 on page 185.

Table 14. Network Attributes

APPN Object Model	APPN Object Model with HPR
<ul style="list-style-type: none"> • Controller description • Device description <ul style="list-style-type: none"> – ROCHESTER – ROCHESTER – DESMOINES01 – STPAUL01 – DESMOINES – ROCHESTER01 – DESMOINES – STPAUL02 – STPAUL – ROCHESTER02 – DESMOINES02 – STPAUL 	<ul style="list-style-type: none"> • Controller description (Real) • There are no device descriptions here • Controller description (Virtual APPN Controller) • QAPEND001 • Device description <ul style="list-style-type: none"> – ROCHESTER – DESMOINES – STPAUL • APPN device descriptions <ul style="list-style-type: none"> – ROCHESTER – DESMOINES – STPAUL

12.2.3 Device Recovery Performance for Display Devices

Pass through or Telnet display devices simulate device recovery for a powered off condition. On systems prior to V4R2, the interactive subsystem attempts to recover the device by sending out a signon screen. This recovery process sends several messages to the job log of the interactive subsystem indicating that an error occurred on the device.

The attempt to send a signon display to the display device is not done on V4R2 systems. This reduces the overhead on the system.

Note: For virtual displays, eliminating this overhead is always good because the recovery attempt *always* fails.

12.2.4 MAXFRAME Value on LAN Controller Descriptions

The Maximum frame size (MAXFRAME) value used for a controller is negotiated at connect time. The actual value used is in message CPF5908 indicating:

```
Controller &1 contacted on line &2...
If controller &1 is on a local area network, the negotiated
frame size is &3.
```

However, since there may be hundreds of messages, it is difficult to find the one needed to determine the actual value used for a particular controller. On V4R2 systems, the active MAXFRAME setting is placed into the controller description. Here, it can be easily viewed using the Display Controller Description (DSPCTLD) command as shown in Figure 55 on page 186.

```

                                Display Controller Description
                                03/20/98 20:27:22 SYSTEMXX
Controller description . . . . . : CTLD      CTLDSP01
Option . . . . . : OPTION    *BASIC
Category of controller . . . . . :           *APPC

Link type . . . . . : LINKTYPE *LAN
Online at IPL . . . . . : ONLINE  *YES
Active switched line . . . . . :           ALINE01
Character code . . . . . : CODE     *EBCDIC
Maximum frame size . . . . . : MAXFRAME 16393
Current maximum frame size . . . . . : 1994
Remote network identifier . . . . . : RMTNETID *NETATR
Remote control point . . . . . : RMTCPNAME SYSTEMAA
Initial connection . . . . . : INLCNN  *DIAL
Dial initiation . . . . . : DIALINIT *LINKTYPE
Switched disconnect . . . . . : SWTDSC  *NO
Data link role . . . . . : ROLE     *NEG

Press Enter to continue.

F3=Exit  F11=Nondisplay keywords  F12=Cancel
More...

```

Figure 55. Current Maximum Frame Size

12.2.5 Force a Vary Off

Use the Vary Configuration (VRYCFG) command force vary off (FRCVRYOFF) parameter to end jobs on APPC devices automatically. This is an alternative solution to responding to CPA2610 inquiry messages on the QSYSOPR message queue indicating:

```
Device &l cannot be varied off at this time ...
There are jobs in the system that are currently using device &l.
```

The intended use for this option is when an error occurs and you cannot vary off a component using normal recovery methods.

Recommendation

Use FRCVRYOFF only when absolutely necessary. The FRCVRYOFF action can leave the associated objects in an unpredictable state. When the jobs end, cleanup does not follow normal paths.

The ability to force a vary off from the VRYCFG command can be used with the FRCVRYOFF parameter to eliminate the inquiry message that occurs when an APPC device has active sessions. This value allows the system to end the associated active jobs automatically.

12.2.6 LAN Response Timer

On APPC, FNC, RWS, RTL and HOST controller descriptions with a link type of *LAN, the value *CALC for the LAN Response Timer (LANRSPTMR) parameter has a default value of 30 (3 seconds). On systems prior to V4R2, the default is 10 (1 second). The increased LANRSPTMR value for V4R2 systems helps prevent unnecessary link-level timeout conditions, which can be common on Ethernet networks with a large amount of activity.

One side-effect of a larger value is that it takes longer for the system to detect a link failure (by default, 30 seconds on V4R2 systems and 10 seconds on prior level systems). Compare the benefits of the time to detect failures to an earlier timeout.

The following example shows how the 10 and 30 seconds are determined:

```
LANRSPTMR X LANFRMRTY = time to detect link failures
  1      X   10      =  10
  3      X   10      =  30
```

12.2.7 Device Allocation Messages

On systems prior to V4R2, QSYSOPR can receive hundreds of device allocated (or cannot allocate device) messages when subsystems are started for devices that are not powered on (or for device descriptions that are defined without associated hardware). The following messages are no longer sent to the QSYSOPR message queue on V4R2 systems:

- CPF1187 Subsystem &1 cannot allocate workstation &2.
- CPF1273 Communications device &2 was allocated to subsystem &1.
- CPF1274 Subsystem &1 cannot allocate communications device &2.
- CPF1275 Subsystem &1 cannot allocate device &2.

The device that is allocated is shown on the Display Device Description display in Figure 56.

```

                                Display Device Description
                                SYSTEMXX
                                01/22/98 20:27:22
Device description . . . . . : DEVD      AAADDEV1
Option . . . . . : OPTION    *BASIC
Category of device . . . . . :          *NET

Device type . . . . . : TYPE      *TCPIP
Online at IPL . . . . . : ONLINE   *NO
Attached controller . . . . . : CTL    SXPCTL
Allocated to:
Job name . . . . . :          QTCPIP
User . . . . . :            QTCP
Number . . . . . :            052651
Text . . . . . : TEXT        CREATED BY AUTO-CONFIGURATION

Press Enter to continue
                                Bottom
```

Figure 56. Active Mode for APPC Devices

Also, the job that an APPC mode is allocated to now appears under Active Mode on the Display Device Description display as shown in the following figure.

```

                                Display Device Description
                                01/22/98 20:22:27
Device description . . . . . : DEVD      QANRDEVA
Option . . . . . : OPTION    *MODE
Category of device . . . . . :          *APPC

-----Active modes-----
Mode      Job      User      Number
QADSM     QCMN     QSYS     052651
SNASVCMG  QLUS     QSYS     020765

Press Enter to continue.
                                Bottom

```

Figure 57. Active Mode for APPC Devices

The messages are still logged to the QHST file, as well as some of the subsystem joblogs.

12.2.8 QAUTOVRT System Value

On systems prior to V4R2, the limit of the number of virtual devices that are autocreated by the system is 9999. As your system and network continue to grow, this can become a limiting factor. On V4R2 systems, the limit is 32500 and *NOMAX.

The names of auto-created virtual devices can include alphanumeric characters as well.

12.2.9 Removal of Obsolete Messages in QSYSOPR

On systems prior to V4R2, when a communications error occurs and automatic recovery attempts are exhausted, an inquiry message is sent to the QSYSOPR message queue. On V4R2 systems, if the associated link recovers from the error without a response to the inquiry message, the unanswered message is removed from the QSYSOPR message queue. One example of a change in state is when the object is varied off. The message is removed when the state of the network interface, line, controller, or device changes. This makes the system operator's job easier. It simulates the appearance of a failed network in QSYSOPR.

Note: There are PTFs available for R410 and R420 systems to allow QSYSOPR to wrap when it fills, instead of issuing a warning message. See 10.6, "QSYSOPR Message Queue Wrap When Full" on page 151, for considerations on when QSYSOPR fills.

12.2.10 TCP Error Recovery

As the fastest growing choice of connection methods for the AS/400 system, TCP has been continually enhanced with changes allowing for easier system administration. This includes:

- Controlling the logging of protocol errors
- Controlling the inactivity timeout for Telnet devices

- Controlling device recovery actions for Telnet devices
- Controlling remote signon actions for Telnet devices

Except for controlling remote signon actions, each enhancement is further described in this section. Refer to *Remote Work Station Support*, SC41-5402, to understand signon actions for remote workstations.

12.2.10.1 Controlling the Logging of Protocol Errors

You can choose to log protocol errors. If the TCP/IP attribute LOGPCLERR is set to *YES, frame and error type logs are written to the system product activity log communications area. As broadcast messages are considered errors, this can be a significant system activity. In networks where hundreds of broadcast messages are received by the AS/400 system, the logging of these broadcast messages can consume 10 to 30% of the CPU. The logging of these messages can also cause the error log data to wrap (logged data is overlaid with new log data), and lose some potential problem determination data. Networks that experience excessive numbers of broadcast messages degrade network availability. In these situations, AS/400 performance is impacted if the TCP/IP attribute LOGPCLERR indicates *YES.

The communications frame and other error information is placed into the communications log portion of the product activity log. In the error log entry formatted in hexadecimal format, the frame contents begin at an offset of X180.

By default on V4R2 systems, broadcast messages are not logged when LOGPCLERR(*YES) is specified. This significantly reduces the impact to the CPU when multiple broadcast messages are received.

On systems prior to V4R1, broadcast message logging can be excluded by setting the TCP/IP packet reassembly time out (IPRSBTIMO) parameter to 79. The number 79 represents a unique timeout value to indicate TCP to exclude broadcast messages.

To get the "old" behavior (prior to V4R2) on V4R2 systems, include broadcast message logging by setting the TCP/IP attribute IPRSBTIMO to 79.

12.2.10.2 Controlling the Inactivity Timeout for Telnet Devices

By popular request and to provide consistency with other communications protocols, beginning with V4R2, Inactive job time-out (QINACTITV) and Inactive job message queue (QINACTMSGQ) system values support Telnet sessions. TCP/IP Telnet, IPX Telnet, workstation gateway, and virtual terminal APIs all support QINACTITV and QINACTMSGQ. PTF SF47141 is available for V4R1 systems to enable this inactivity timeout support for Telnet devices. See the cover letter for instructions.

Recommendation

If you used the INACTTIMO parameter on the Change TELNET Attributes (CHGTELNA) or Change Work Station Gateway Attributes (CHGWSGA) commands, set the INACTTIMO value to 0 (no timeout) and use QINACTITV instead.

For QINACTITV and INACTTIMO, whichever timer expires first is honored.

12.2.10.3 Controlling Device Recovery Actions for Telnet Devices

On V4R2 systems, the device recovery action (DEVRCYACN) parameter options—disconnect end request (*DSCENDRQS) and disconnect message (*DSCMSG), specifically support Telnet devices. This allows Telnet sessions that are impacted by a communications failure to be disconnected, rather than ended. The disconnected job is reconnected when the user recovers the Telnet session.

Note: You must use the Telnet exit programs, and specify a device name (non_QPADEV*) to enable this function. Or, you can:

- Use network stations.
- Use the latest Client Access/400 V3R1M3 service pack.

Refer to APAR II10918 for considerations. See *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01, for more information on accessing APARs.

12.2.11 Serviceability Improvements

Service personnel, including the system administrator, use diagnostic tools provided with OS/400 to help resolve problems and maintain system availability. The tools are used to:

- Analyze the Vertical Licensed Internal Code Logs (VLOGs)
- Run Licensed Internal Code (LIC) traces
- Run communications traces

The enhancements to each of these tools are described further in this section.

12.2.11.1 VLOG Improvements

Work has been done to limit the potential to produce many duplicate VLOGs for APPC communications. If a problem condition occurs so that many configuration objects fail on systems prior to V4R2, the system generates a VLOG entry for every affected object. Plus, a log entry is written to the QSYSOPR message queue for every affected configuration object. This overhead is reduced on V4R2 systems. Generally, in V4R2, only the highest-level object in the hierarchy that is affected generates a VLOG and posts a message to the QSYSOPR message queue.

For example, if a defect condition occurs on a line, on systems prior to V4R2, there are VLOGs produced for the line, all attached controllers, and all controllers' attached devices. Likewise, messages indicating the object's condition as unusable are sent to the QSYSOPR message queue for every object. On V4R2 systems, there is only one VLOG for the line failure and one message to the QSYSOPR message queue indicating that the line is unusable.

This saves on system ERP overhead and makes the system easier to administer with fewer logs and messages to manage.

Note: Use the WRKCFGSTS command for the object indicated in the message posted to QSYSOPR to see all the affected objects. Then, vary them off.

12.2.11.2 Licensed Internal Code Trace Improvements

How a LIC trace is completed for the components that support APPC communications is improved to limit the amount of data traced. The System Service Tools (SST) LIC interface for a Source Sink Component trace on V4R2 systems makes setting up LIC traces more efficient. You can request to trace all objects attached to a higher level communication configuration object, rather than step through each object separately. For example, you can define a trace for a line and specific controllers. You only need to define the trace once. It can run as frequently as required.

You might consider pre-defining a trace for those combinations of objects used most often.

12.2.11.3 Communications Trace Improvements

The communications trace maximum allowed buffer size is up to 64MB on V4R2 systems. In addition, a filtering capability is available for Token-Ring lines running on the Token-Ring IOPs, supported with V4R2 hardware.

This allows for more data to be captured for communication failures, which is a benefit when errors are difficult to time or trap.

Note: Refer to the *AS/400 Advanced Series System Handbook*, GA19-5486, for a list of Token-Ring IOPs supported for each operating system release.

12.3 Improvements on V4R1 Systems

Since V4R1, the focus has been to prevent excessive recovery times and reduce required system resources. This allows for faster connection, disconnection, and reconnection times so that end users can resume productivity more quickly after a failure.

Communications error recovery procedure improvements in V4R1 were targeted for networks running APPC. Users with PCs using Client Access/400 over APPC observe the most benefits.

When system resources are reduced, performance is improved. V4R1 and later systems offer reduced resource usage with:

- A restructure of the target pass-through job
- A restructure of the file-server jobs
- An optional activation of the LAN Manager
- Limits to the occurrence of program start request (PSR) messages (CPF1269) with Reason Code 401 (which usually indicates the PSR was received for a device not allocated to an active subsystem)
- Improvements to the filtering of error-log messages
- Changes to the communications trace function

These improvements are discussed further in the following sections.

12.3.1 Restructure of 5250 Display Station Pass Through

Target 5250 display station pass through is used for:

- System-to-AS/400 signon sessions

The session can be another AS/400 system or another system supporting pass through (for example, a System/36).

- APPC-based PC emulators, such as PCOMM, PC5250, and others
This case is also known as the Work Station Function for Client Access/400.
- Any other 5250 support written using the pass-through Licensed Internal Program Interface (LIPI)

The 5250 display station pass through (also known as target work station function) replaces communication jobs with a set of server jobs. Changes visible to the end user include:

- The QPASTHRSVR system value
- No jobs in the QCMN subsystem for 5250 target pass-through sessions (removes the performance impact for job-start and job-end processing required for every 5250 target display station pass-through session)
- One primary server job and “n” secondary server jobs in the QSYSWRK subsystem

On systems prior to V4R1, every 5250 signon session that uses 5250 target display station pass through has two jobs associated with it:

- The interactive user job, which is typically in the QINTER subsystem
- The communications job, which is typically in the QCMN subsystem

The communications job is responsible for activation and deactivation processing for the pass-through session.

An example scenario showing the affect of these changes is:

Consider a system that has 300 PC users, each with a 5250 session to the AS/400. This results in a total of 600 jobs on the AS/400, all of which have to be ended in order to recover from a network communication failure.

On V4R1 systems, there is one interactive job for each 5250 session. With our scenario, there are 300 jobs to be ended in order to recover from a network failure rather than 600.

The target display status for pass-through jobs is identified on the WRKCFGSTS display with a job name of QPASVRP/QSYS/number as shown in Figure 58 on page 193 and Figure 59 on page 193.

```

                                Work with Configuration Status
                                02/08/98 13:02:39 SYSTEMXX
Type options, press Enter.
  1=Vary on  2=Vary off    5=Work with  8=Work with description
  9=Display mode status ...

Opt Description Status                -----Job-----

SYSTEMNN      ACTIVE
SYSTEMNN      ACTIVE
BLANK         ACTIVE/TARGET      QPASVRP  QSYS  00289
SYSTEMMM      ACTIVE
SYSTEMMM      ACTIVE
BLANK         ACTIVE/TARGET      QPASVRP  QSYS  00289
SYSTEMPP      VARY ON PENDING
SYSTEMPP      VARIED OFF
BLANK         VARY ON PENDING

Parameters or command
===>
F3=Exit  F4=Prompt  F12=Cancel  F23=More options  F24=More keys

```

Figure 58. QPASVRP Job on WRKCFGSTS Screen

```

                                Work with Configuration Status
                                05/04/98 11:37:40 SYSTEMXX
Position to . . . . . _____ Starting characters

Type options, press Enter.
  1=Vary on  2=Vary off    5=Work with job  8=Work with description
  9=Display mode status  13=Work with APPN status...

Opt Description Status                -----Job-----

--- SYSTEMZZ      VARIED OFF
--- SYSTEM05      VARIED OFF
--- SYSTEM0500    VARIED OFF
--- SYSTEMYYSP    VARIED OFF
--- SYSTEMAA      ACTIVE
--- SYSTEMAA      ACTIVE
--- BLANK         ACTIVE/TARGET      *PASSTHR
--- BLANK         ACTIVE/SOURCE      QPADEV000I  DMTSCP  205700
--- SYSTEM26      VARY ON PENDING

More...

Parameters or command
===>
F3=Exit  F4=Prompt  F12=Cancel  F23=More options  F24=More keys

```

Figure 59. WRKCFGSTS Screen After PTFs Applied

12.3.1.1 A Change to the WRKCFGSTS Display

With V4R1, pass-through server jobs replaced the communications jobs. As a result, the primary pass-through server job (named QPASVRP QSYS xxxxxx) appears on the WRKCFGSTS display for the APPC devices as the job for all target pass-through sessions (including the target job for PC Workstation Function, PCOM, and so on). When PCs hang or otherwise become unusable, it is a common occurrence for operators to end all jobs associated with the PC's APPC device. Thus, the QPASVRP job is ended, which takes down all of the current pass-through users.

To make it more difficult for someone to cancel the primary pass-through server job, a change has been made in the information shown on the WRKCFGSTS display for the APPC device description. When a session is for target pass-through, the display shows the value, *PASSTHR, instead of the primary pass-through server jobname. Also, the job options for the modes with *PASSTHR do not work. The WRKCFGSTS appearance helps to "hide" the job from the operator.

Operators are forced to vary off the APPC device while it is use, rather than ending all jobs on the APPC device before varying it off.

Note: It is still possible to do an ENDJOB on the primary pass-through server job. However, this change helps "hide" the job from the operator.

This change is described in APAR SA72607. Refer to *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01, for information on accessing APARS.

Without the PTF, operators can "mistakenly" end all jobs on an APPC device before varying the device off. To vary the APPC device off on a system with the PTF applied, the user must perform one of the following options:

- Option 1
 1. Find the virtual device description for the user.
 2. Issue an ENDJOB for each device that has a job.
 3. Vary off the associated virtual device.
- Option 2
 1. Vary off the APPC device while it has active sessions.
 2. Respond to inquiry message CPA2610 in the QSYSOPR message queue with a G *or*
 3. On V4R2 systems, use the VRYCFG command with the FRCVRYOFF parameter.

12.3.1.2 QPASTHRSVR System Value

The QPASTHRSVR system value specifies the number of target display station pass through server jobs that are available. These jobs process AS/400 display station pass through, AS/400 Client Access Work Station Function (WSF), and other 5250 emulation programs on programmable workstations.

The operating system calculates a reasonable number of target display station pass through server jobs. Four secondary server jobs for each CPU are started when using the default setting.

Recommendation

A "0" value for the number of pass-through servers is used for migration during an initial installation only. Do not use a "0" value after the migration is successfully completed. Support for a value of "0" may be dropped in future releases beyond V4R2.

Note: The server jobs are not used for Telnet and virtual terminal manager (VTM) devices. Therefore, if you use only Telnet or VTM, you may want to decrease the value specified for the number of target display station pass-through server jobs. The following display highlights the jobs started when the QPASTHRSVR system value is a non-zero value.

```
Work with Active Jobs                                     SYSTEMXX
                                                         12/08/97 13:02:39
CPU %:      .0      Elapsed time: 00:00:00      Active jobs: 211

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      Type  CPU %  Function      Status
---
   QSYSWRK      QSYS      SBS    .0
   ADMIN        QTMHHTTP  BCH    .0  PGM-QTMHHTTP  TIMW
   ADMIN        QTMHHTTP  BCI    .0
   QNSCRMON     QSVSM     BCH    .0  PGM-QNSCRMON  DEQW
   .            .         .      .
   .            .         .      .
   .            .         .      .
   QPASVRP     QSYS      BCH    .0  PGM-QPASVRP   DEQW
   QPASVRS     QSYS      BCH    .0  PGM-QPASVRS   TIMW
   QPASVRS     QSYS      BCH    .0  PGM-QPASVRS   TIMW
   QPASVRS     QSYS      BCH    .0  PGM-QPASVRS   TIMW
   QPASVRS     QSYS      BCH    .0  PGM-QPASVRS   TIMW
   .            .         .      .
   .            .         .      .
   .            .         .      .
   QPRFSYNCH   QSYS      BCH    .0  PGM-QFPAPRFJ  DEQW
More...

Parameters or command
====>
F3=Exit  F4=Prompt      F5=Refresh  F10=Restart statistics
F11=Display elapsed data  F12=Cancel  F14=Include  F24=More keys
```

Figure 60. Pass-Through Server Jobs

The following sections describe the QPASVRP and QPASVRS server jobs.

12.3.1.3 Pass-Through Server Jobs

There are several autostarted server jobs in the QSYSWRK subsystem:

- One primary server job named QPASVRP/QSYS/nnnnnn to manage the secondary server jobs
- Below QPASVRP, between 1 and 100 QPASVRS jobs named QPASVRS/QSYS/nnnnnn to perform the work previously done by the communication jobs

Multiple server jobs enable a performance improvement for target pass-through startup, termination, and error recovery. On systems prior to V4R1, there may be hundreds of jobs in the QCMN subsystem for target pass through. On V4R1 systems, there is by default one job in the QCMN subsystem, plus four server jobs for each CPU.

For example, on a 12-processor system, there are 49 jobs calculated using the following equation:

$$12 \text{ CPUs} \times 4 \text{ server jobs per CPU} = 48 + 1 = 49 \text{ jobs}$$

For more information, see *Remote Work Station Support*, SC41-5402.

Note: Job start and termination on the AS/400 system is resource intensive. By reducing the amount of communications jobs used for target pass through (that *only* performs activation and termination functions) the number of job starts and ends is reduced.

The positive effect of server job implementation is in terms of the starting and ending pass-through sessions:

- Clients connect faster when users begin their day.
- Clients disconnect faster when users end their day.
- Clients recover faster from network outages or network failures.

This support applies to SNA STRPASTHR and 5250 emulation within an APPC data stream, such as with PC5250, Personal Communications, and OS/2 Communications Manager.

Note: For V4R1 and V4R2, this implementation does not apply to Telnet 5250 sessions.

Figure 61 and Figure 62 show the differences in the WRKACTJOB command output on systems prior to V4R1, compared to V4R1 and later.

Subsystem/Job	User	Type	CPU %	Function	Status
QCMN	QSYS	SBS	.0		DEQW
P23CDR75	A97011862	EVK	.0	* -PASSTHRU	EVTW
P23CFC43	A97011863	EVK	.0	* -PASSTHRU	EVTW
P23CNL89	USERXX	EVK	.0	* -PASSTHRU	EVTW
P97V6PK7	A97011861	EVK	.0	* -PASSTHRU	EVTW

Figure 61. WRKACTJOB with Pass Through Jobs Prior to V4R1

Subsystem/Job	User	Type	CPU %	Function	Status
QSYSWRK	QSYS	SBS	.0		DEQW
QPASVRP	QSYS	BCH	.0	PGM-QPASVRP	DEQW
QPASVRS	QSYS	BCH	.0	PGM-QPASVRS	TIMW
QPASVRS	QSYS	BCH	.0	PGM-QPASVRS	TIMW
QPASVRS	QSYS	BCH	.0	PGM-QPASVRS	TIMW
:	:			:	:

Figure 62. WRKACTJOB SBS(QSYSWRK) on V4R1

Note:

- There are no communications jobs for the pass-through sessions in the QCMN subsystem.
- If QPASTHRSVR is zero, pass-through jobs are routed as before to the QCMN subsystem (by default).

Recommendation

Since there are no longer jobs in the communications subsystem (usually QCMN) for pass through, communications entries are no longer used. If you rely on communications entries for security, you can no longer do this the same way when using pass-through servers. You can, however, use the QRMTSIGN system value program to designate how you want to allow remote signon.

For further information on controlling remote signons, refer to *Tips and Tools for Securing Your AS/400*, SC41-5300-01.

Recommendation

Add applicable routing entries to the communications subsystem affected to control security.

If QPASTHRSVR is set to a non-zero value, pass-through and file server requests do not go through QCMN. Communications entries that *were* used for security purposes for these applications include routing entries. Since these jobs no longer pass through QCMN, the planned for security check does not occur. The routing entries need to be moved appropriately. Be aware that other program start requests *will* use QCMN so some of the communication and routing entries may still apply.

12.3.2 File Server Job Restructure

The file server (shared folders) start-up and take-down paths were optimized on V4R1 and later systems. Changes visible to the end user include:

- The QSERVER subsystem is started at IPL.
The system supplied start-up program starts the QSERVER subsystem. On systems prior to V4R1, the system checks to see if QSERVER needs to be started when a file server program start request is received.
- If QSERVER is not started, the file server program start requests are rejected with a message CPF1269 in QSYSOPR indicating:
Program start request received on communications device was rejected with reason code 413
No file serving is available.
- The program start request goes directly to the QSERVER subsystem. The file server no longer has a job that starts or ends in the QCMN subsystem. As with pass-through changes, the server does not go through the QCMN subsystem. This means that the communications entries in QCMN are not effective for the file server.
- All file server jobs are now pre-start jobs in the QSERVER subsystem to handle network shared folders work.

- The file server performs a faster start-up due to fewer jobs and better repeat use of the prestart jobs.
- On systems prior to V4R1, there are two job starts for every client. The communications job in QCMN followed by the job in QSERVER jobs need to communicate. The elimination of the communications job reduces work on the system and simplifies the start-up and take down logic, as well as error recovery.
- Using WRKCFGSTS on the APPC controller description shows the file-server jobs. These file-server jobs have message CPIAD12 logged to the job log indicating:

```

Servicing user profile &1 from client &2...
This prestart job is currently servicing requests for
profile &1 from client &2. The client name is either an APPC remote
location name, or a TCP/IP remote system name.

```

Note that the message text contains the user profile and client names for which the job is providing service.

12.3.3 Activation of the Operational LAN Manager

Availability of a LAN network depends on having tools available to access the pertinent information for failure diagnosis. On a LAN line description, the administrator can choose to activate the LAN manager function by selecting *YES for the ACTLANMGR parameter when the LAN line is created or changed. The LAN manager function is activated when the line is varied on. Likewise, to deactivate the LAN manager, the Token-Ring line must be varied off, and back on again.

If ACTLANMGR(*NO) is selected, the ring-error monitor (REM) and configuration ring services (CRS) functions are not activated in the IOP. Plus, configuration changes are not logged and beacon errors are not posted. If ACTLANMGR(*NO) is specified, less disk space and disk arm activity is required. This option also offers the following advantages:

- Less code is required in the IOP.
- It eliminates the posting of LAN Manager messages to QSYSOPR and QHST.
- It eliminates a duplication of work between the LIC and the QSYSCOMM1 system job.

On busy networks the ACTLANMGR feature is usually turned off. However, consider carefully the usefulness of the error log information before deactivating this feature.

12.3.4 Program Start Request Message Threshold

The error message CPF1269 with reason code 401 is received when the device is not allocated to an active subsystem (usually QCMN). This is very common while the system is in a restricted state. This message states:

```

Program Start Request received on device not allocated to active
subsystem message

```

Historically, these CPF1269 reason code 401 errors have been known to flood the QSYSOPR message queue (or QSYSMSG if it exists) for each active device and mode pair. Since PCs typically attempt to connect to the system repeatedly until successful, this problem is compounded on systems with many LAN attached PCs.

The number of times that a CPF1269 message with a reason code 401 error is sent for a device and mode combination is limited to one. This change also applies to cases where message CPF1269 with a reason code 401 is caused by a configuration error (that is, there is no subsystem configured for the device). Varying the affected device on and off again causes a repeat of the CPF1269 message.

Note: Messages other than CPF1269 with reason codes other than 401 are not affected.

12.3.5 Error Log Filtering

Historically, problem log errors can flood the system causing an excessive number of:

- Product activity logs
- Problem logs
- Messages to the QSYSOPR message queue

Overall, system performance is degraded noticeably, as the QSYSCOMM1 system job is burdened with producing error logs. On V4R1 systems, error analysis support in LIC includes a filter that takes affect after an initial burst of 100 identical errors. It is limited to 100 identical errors per hour thereafter.

12.3.6 Communications Trace Improvements

As systems and networks continue to grow in capacity, serviceability and debug functions become more critical. Communications tracing tool capabilities are enhanced to include:

- A buffer size increased to 16MB to capture errors more quickly
- Options to select remote controllers, IP address, IP protocol, and the local SAP or remote SAP number (with V4Rn IOPs only) when starting a communications trace
- Communications trace time stamps no longer need to be adjusted by the user to match time stamps provided by the hardware (with V4Rn IOPs only)

The Record Timer column in the communications trace report is presented in one of two formats described in the field description in Figure 63 on page 200.

```

                                Display Spooled File                                SYSTEMXX
                                                                                   02/27/98   13:02:39
File . . . . . : QPCSMPT
Control . . . .
Find . . . . .
COMMUNICATIONS TRACE      Title: TRACE                                02/27/98 10:14:00
Record Number . . . . . Number of record in trace buffer (decimal)
S/R . . . . . S=Sent R=Received M=Modem Change
Data Length . . . . . Amount of data in record (decimal)
Record Status . . . . . Status of record
Record Timer . . . . . Time stamp. Based on communications hardware, the time
                        stamp will be either:
                        1. 10 microsecond resolution time of day
                           (HH:MM:SS.NNNNN) based on the system time
                           the trace was stopped
                        2. 100 millisecond resolution relative timer with
                           decimal times ranging from 0 to 6553.5 seconds
Data Type . . . . . EBCDIC data, ASCII data or Blank=Unknown
Controller name . . . . Name of controller associated with record
Command . . . . . Command/Response information
Number sent . . . . . Count of records sent
Number received . . . . Count of records

```

Figure 63. Record Timer on a Communications Trace

12.4 Configuration Tips and Techniques

As seen in several sections in this chapter, how your system is configured can make a significant difference in its performance during error recovery. Further configuration considerations that can have a significant impact on APPC communications error recovery performance include:

- Subsystem configuration
- How the ONLINE parameter is set
- How the SWTDSC parameter is set
- Minimum switched status for APPN
- Communications recovery limits
- Whether APPN controller descriptions are allowed to autcreate
- Whether automatic device deletion is activated
- How prestart jobs are configured
- Client Access mode description management
- Job log considerations

These each are further described in the sections that follow.

12.4.1 Subsystem Configuration

As the number of users on the system increases, it becomes more important to consider how the communications and interactive subsystems are configured.

Interactive subsystems perform device recovery when 5250 sessions end, whether normally or abnormally. During error recovery, when many users lose their sessions at once, an interactive subsystem can become very busy performing device recovery. In addition, this device recovery can adversely impact the work of other users in the subsystem that otherwise may be unaffected by the failure.

The configuration of subsystems has little impact for normal data path operations. However, multiple subsystems can provide multiple processes to perform cleanup and recovery when an error condition occurs. This can result in significantly improved recovery performance.

You can also group devices into subsystems by specifying those devices *not* to allocate to the subsystem. To do this, run the Add Work Station Entry (ADDWSE) command:

```
ADDWSE SBSDB(library-name/subsystem-name)
WRKSTN(device-name) AT(*ENTER)
```

You can execute the ADDWSE command while the subsystem is active. However, subsystems do not reallocate devices dynamically. It may be necessary to end and restart the subsystem to have the device allocated as desired.

Similar considerations apply to communications subsystems. Use the ADDCMNE command to set up the appropriate communications entries.

Recommendation

Use a soft limit for the number of devices serviced by a single subsystem. A configuration with 200 to 300 devices per interactive subsystem is a reasonable goal. Create additional communications and interactive subsystems to split the work into multiple subsystems if necessary.

Note: Do *not* add entries for all other devices at *ENTER. *ENTER devices consume a lot of system resources when subsystems are started and ended when there are many devices running over a switched line.

When dividing work into multiple subsystems, consider the following factors:

- The number of users in any given subsystem
- The connectivity used to access the system

You may want all of the users on one Token-Ring LAN to run in one subsystem. Those users connected using a 5494 remote controller can benefit by running in a different subsystem.

- The type of work the users do
- The geographic location of the users

12.4.2 Online at IPL Considerations

The online at IPL parameter and controller configuration objects are set by default to *YES. This may not be a good choice for all environments. In particular, we recommend that you do not use the ONLINE(*YES) option for configuration objects that end up in recovery when they are varied on. This results in unnecessary work (communication recovery) during the IPL.

It can be difficult to predict which objects do not vary on successfully. APPC controller descriptions that indicate *DIAL out to PCs are prime examples.

Consider not varying on the objects at IPL and manage the varying on of these objects yourself with a CL program or the system start-up program. For controllers on LANs, use the auto-configuration parameter (AUTOVRTCTL(*YES))

on the appropriate LAN line description and let the system vary the object on when first used.

Recommendation

Avoid varying on network server descriptions during an IPL. When network server descriptions are varied on, the integrated PC server is reset. This reset process can take a long time. During the IPL, this work is done by the QSYSCOMM1 arbiter job, which means QSYSCOMM1 is less available to do other work.

Refer to *Communications Management*, SC41-5406; *Communications Configuration*, SC41-5401; and *APPN Support*, SC41-5407, for further information about the online at IPL parameter.

12.4.3 Switched Disconnect

For many switched environments, the connection should be disconnected when a user logs off the host. This is especially beneficial when long distance charges incur during the connection. However, for LAN attached PCs, this may is not the best configuration.

Although the PC connects to the AS/400 system, the application may not start soon enough, or the router starts, but not the 5250 session. In this case, the connection to the AS/400 system may be automatically disconnected.

Recommendation

Change the default of the switched disconnect (SWTDSC) parameter to *NO for LAN attached devices.

Using SWTDSC(*NO) keeps the connection to the AS/400 system active even if no applications are active.

The SWTDSC parameter is discussed further in the *Communications Configuration*, SC41-5401, and *Communications Management*, SC41-5406.

12.4.4 APPN Minimum Switched Status

In an APPN network, a minimum switched status (MINSWTSTS) of *VRYON can limit the routes that APPN thinks are available. This prevents APPN from choosing routes that have a controller in a varied on pending status that are on the system, but are varied off or inoperative on the adjacent system.

The default value for MINSWTSTS is *VRYONPND. It allows APPN controllers in a vary on pending status to be available for APPN route selection. This may or may not be appropriate for your network.

Note: MINSWTSTS(VRYON) requires SWTDSC(*NO).

For more information on the APPN minimum switched status parameter, see *Communications Configuration*, SC41-5401, and *APPN Support*, SC41-5407.

12.4.5 Communications Recovery

Before automatic communications recovery is involved, first-level recovery is performed. Each link protocol defines the timers and retry values that apply to first-level recovery. The system and applications are unaware of the low-level recovery that occurs. Excessive recovery appears as poor performance.

Second-level recovery of communications devices consumes resources. The system handles the error and may attempt to reconnect to the remote system. If, however, the attempt to reconnect fails (as is generally the case with PCs on a LAN), disabling second level recovery reduces the system resources that are used.

The state in which the configuration object is left after a failure differs depending on whether second level recovery is on or off. Using second level recovery, the object is in a varied on pending (VRYPND) status and a message is posted to the QSYSOPR message queue. If second level recovery is disabled, the same object is in a recovery-pending status (RCYPND). A message is also posted to the QSYSOPR message queue.

There is a trade-off between consuming more CPU resource for automatic error recovery and requiring a more manual intervention for recovery.

Note: The first recovery is still done regardless of whether the recovery limit is set to 0. On a LAN, this means that the inactivity timer parameter (INACTTMR) is used to determine if the remote system is still available. Once the inactivity timer expires, first level recovery is driven by the LANFRMRTY and LANRSPTMR parameters. A suggestion is to change the LANRSPTMR value on the controller description to 30 (3 seconds). Then, with a LANFRMRTY value of 10, the system is allowed 30 seconds to determine if the remote connection is still available.

Note: Turning off second level recovery can cause the devices/controllers to go into a RCYPND status and require operator intervention. A value of (0 5) on the communications recovery limit parameter (CMNRCYLMT) disables communications error recovery. The value of (2 5) is the default—which requests a reconnection twice within five minutes.

Due to the relationship between the retry count and the time interval, in cases where the *additional retries* value is greater than the *time interval* value it is possible that the system remains in automatic recovery if the time interval expires before the retry limit is reached.

Recommendation

Keep the additional retries value less than the time interval value on the QCMNRCYLMT system value as well as on the controller and object description CMNRCYLMT parameter.

For additional information, refer to the INACTTMR parameter in the *Communication Configuration*, SC41-5401.

Also note that manual recovery is not the only option. Applications can be written to determine if a failure has occurred and handle it accordingly. If you write such applications, consider the following points:

- Monitor for error messages in QSYSOPR. When they occur, handle the condition.

- Monitor the status of configuration objects. If the object status is not active, varied on, or varied on pending, take the appropriate action.

Some useful APIs for checking the status of configuration objects include:

- Retrieve Configuration Descriptions (QDCRCFGS)
- List Configuration Descriptions (QDCLCFGD)

12.4.6 APPC Controller Description Error Recovery

The interaction between many of the parameters that affect the system action taken when APPC controller descriptions enter recovery are shown in Table 15 and Table 16.

MINSWTSTS	INCLNN	APPN	CTLOWN	Power PC Off (Recovery)	Manual Vary On
*VRYONPND	*DIAL	*YES	*SYS	Dial not attempted	Dial is attempted
*VRYONPND	*DIAL	*YES	*USER	Dial is attempted	Dial is attempted
*VRYONPND	*DIAL	*NO	*SYS	Configuration not allowed	Configuration not allowed
*VRYONPND	*DIAL	*NO	*USER	Dial not attempted	Dial is attempted
*VRYONPND	*ANS	*YES	*SYS	Dial not attempted	Dial not attempted
*VRYONPND	*ANS	*YES	*USER	Dial not attempted	Dial not attempted
*VRYONPND	*ANS	*NO	*SYS	Configuration not allowed	Configuration not allowed
*VRYONPND	*ANS	*NO	*USER	Dial not attempted	Dial not attempted

Note: Environments other than APPC-attached PCs on a LAN using Client Access/400 may result in different behaviors.

APPN	INLCNN	CTLOWN	SWTDSC	Power PC Off (Recovery)	Manual Vary On
*YES	*DIAL	*SYS	*YES	Configuration not allowed	Configuration not allowed
*YES	*DIAL	*SYS	*NO	Dial is attempted	Dial is attempted
*YES	*DIAL	*USER	*YES	Configuration not allowed	Configuration not allowed
*YES	*DIAL	*USER	*NO	Dial is attempted	Dial is attempted

Note: The tables relate to APPC-attached PCs on a LAN using Client Access/400 and assume the following factors:

1. DIALINIT(*IMMED) is specified on the controller description.

2. There are no APPN CP sessions.

Other references on this topic can be found in the *Communications Configuration*, SC41-5401; *Communications Management*, SC41-5406; and redbook *IBM AS/400 V3 Communication API Handbook*, SG24-2573.

12.4.7 Automatic Creation of APPC Controllers

On systems prior to V4R2, automatic creation of APPC controllers and device descriptions is handled by the QLUUS system job. On V4R2 systems, the same function is done in the QCMNARBnn system jobs. Automatically created APPC controllers are created by default with the following parameters:

- ONLINE(*NO)
- INLCNN(*DIAL)
- DIALINIT(*LINKTYPE)
- APPN(*YES)
- SWTDSC(*YES)
- MINSWTSTS(*VRYONPND)
- AUTODLTDEV(1440)

These values were selected to provide the best overall solution for both PCs and system-to-system communications. However, they may not meet the specific requirements for your system.

Recommendation

If you find unnecessary recovery attempts using the default values, consider using a model controller and modify the parameter values as suggested in this chapter. A model controller is necessary because you cannot change the command defaults since APPN automatic configuration does not use defaults. Code all the command parameters.

Refer to *Communications Configuration*, SC41-5401, and *APPN Support*, SC41-5407, for further information on model controllers.

12.4.8 Automatic Deletion of Controllers and Devices

Similar to the autocreate function of APPC controllers and devices, the autodelete function is done in the QLUUS system job on systems prior to V4R2. On V4R2 systems, the QCMNARBnn system jobs perform this work. The default behavior is set to automatically delete controllers and devices that are automatically created. This function does *not* automatically delete controller or devices that are created **manually**. Therefore, to prevent devices from being deleted automatically, change the device description. The object is then considered to be created manually.

Note: You cannot turn off the automatic creation of APPN devices.

The default time interval in which to delete APPC controllers (AUTODLTCTL parameter on the line description) is 1 440 minutes (24 hours). This default may not be appropriate for many users. For example, when business resumes on Monday morning, more than 24 hours has passed. Therefore, objects that will be used on Monday morning are deleted and created again automatically upon first access. This causes an unnecessary workload, especially when users attempt connection at roughly the same time.

Recommendation

Remember to set the interval accordingly for weekends and holidays, and set the AUTODLTCTL value for the most common working environment.

Recommendation

APPN networks with multiple routes through the network can result in multiple device descriptions. On V4R2 systems, you can manage this by using the allow APPN virtual support (ALWVRTAPPN) network attribute to use virtual APPN controllers.

For more information on virtual controller considerations, refer to *APPN Support*, SC41-5407.

12.4.9 Prestart Jobs

Prestart jobs are jobs that start when the subsystem starts. The jobs are reused rather than started and ended repeatedly. This is more important for short transactions, where the start and end times comprise a larger percentage of the total time for the job. Prestart jobs allow users to reconnect more quickly after an error.

To use prestart jobs, code a loop program to accept and process the incoming program-start request in the user application. Loop to process a new request, and if an error occurs, end the job so debug information can be gathered.

Prestart jobs are used for many APPC base host server functions, such as:

- File Server—QPWFSESRV in QSYS
- Central Server—QZSCSRVR in QIWS
- APPC Signon Server program—QACSOTP in QSYS
- Remote Command Server—QZRCSRVR in QIWS
- License Management Server—QLZPSERV in QIWS

Prestart job entries for these server jobs are supplied with the system for the QCMN, QBASE, and QSERVER subsystems. Modify the prestart job entries appropriately taking into consideration:

- STRJOBS(*YES/*NO)
- INLJOBS to set the number of available jobs
Set to a larger number if you have many users to connect to the system, and connect processing needs to be done as quickly as possible.
- THRESHOLD and ADLJOBS to set the threshold and additional jobs values
The THRESHOLD value may be below the total number of active users and ADLJOBS may indicate more jobs than will ever be used. In this situation, jobs unnecessarily start and end without being used, causing unnecessary overhead. It is best to spread job starts out over time.
- MAXJOBS

An example of a prestart job entry set up is:

```
CHGPJE SBS(D) QCMN PGM(QACSOTP) STRJOBS(*YES) INLJOBS(10)
      THRESHOLD(2) ADLJOBS(10) MAXJOBS(*NOMAX)
```


When the QCMN subsystem is started, ten jobs for APPC signon processing are started. When two of these jobs remain, another ten are started. If ADLJOBS is set too large, the system spends too much time starting the additional prestart jobs at the same time.

Furthermore, as user applications are developed, consider using prestart jobs to reduce program start request start-up processing. Refer to *Work Management Guide*, SC41-5306; *APPC Programming*, SC41-5443; and *AS/400 Client Access Host Servers*, SC41-5740, for further information on communication prestart jobs.

12.4.10 Client Access Mode Descriptions

If you use Client Access/400, QPCSUPP, and QSERVER modes are required to be available. If the modes are deleted, error messages in the QSYSOPR message queue indicate the modes are missing on the source AS/400 system, with a CPI5970 message which indicates:

```
session maximum not established for mode &2 device &1
```

On the target AS/400 system, it issues the message CPF594C, which indicates:
Mode &1 specified for remote location &2 is not configured

Any new attempts to connect with Client Access fail until the mode is recreated. **Do not delete these modes!** The typical scenario is for users to delete the configuration object when displaying the status from the WRKCFGSTS display—a common request, but a dangerous result.

12.4.11 Job Log Considerations

Job log generation can impact system performance during error recovery, especially when many jobs end at one time. You may consider not generating job logs.

To specify not to generate job logs, perform either of the following options:

- Change the device recovery action parameter to *ENDJOBNO LIST in the job description or use the CHGJOB command. Note that there is also a QDEVRCYACN system value for ease in configuration if the majority of the job descriptions need to be changed.
- Change the job description *message logging* value to *LOG(4 0 *NOLIST)*.

In this case, job logs are not generated if ended normally, but are generated if jobs end abnormally.

DEVRCYACN also supports options to disconnect jobs rather than end them.

Note: Disconnected jobs consume resources. For example, the system Work Control Block Table (WCBT) can grow. This can have other side effects on performance. Therefore, it is not good to disconnect from jobs to which you will never reconnect.

On the other hand, if you do not have users that will reconnect after a failure, the disconnect options offer improved performance.

Recommendation

Use disconnected jobs only if you know that the users plan to reconnect. This avoids orphaned jobs to which you cannot reconnect.

The system value QDSCJOBTV controls when orphaned disconnected jobs are cleaned up.

Note: The *Communications Management* manual, SC41-5406, describes the techniques mentioned in the chapter titled “Handling Communication Errors.” Refer to 10.2, “Work Control Block Table” on page 141 for a discussion of managing the WCBT.

12.5 Testing Error Recovery

The following tips are geared toward testing the affect of communications error recovery on the application environment. We recommend that senior network support staff who have an understanding of the impact of testing perform these tests.

12.5.1 Types of Error Recovery Testing

There are three basic types of error recovery failures:

- Network
- Client
- Server

Each has unique symptoms and recovery procedures. These areas are further described in this section.

12.5.1.1 Network Failures

You can simulate many network failures by disconnecting cables attached to communications components. Errors seen on the AS/400 system when a cable is disconnected depends on where the disconnection occurs—at the AS/400 end or within the network. If the cable is pulled from the adapter card, the AS/400 system reports a line failure. If the cable is pulled from a hub, multistation access unit (MAU), or somewhere else within the network, the AS/400 system detects the failure due to timeout conditions that occur. Disconnecting cables between networked AS/400 systems results in a line failure on one AS/400 system. The other AS/400 system sees the behavior as if a router or bridge failed within the network.

When disconnecting cables, you need to consider how long to leave the cable unplugged. Two considerations are:

- Unplug the cable, and plug it back in after a *short* amount of time.
Short is the period of time during which the IOP can handle a brief disconnection and be reconnected before higher levels of software become aware that anything happened. *Short* depends on the protocol being used. For Token-Ring networks, this is about 60 seconds.
- Unplug the cable and leave it unplugged. This simulates a line failure on the local system and a timeout condition on the remote system.

It is difficult to force a router or bridge failure because this impacts too many users outside the AS/400 network. Typically, disconnecting a cable can simulate a router failure on one side and simulate a line failure on the other side.

12.5.1.2 Client Failures

Commonly used clients are PCs and network stations. Failures that can be forced on the client side of the connection include:

- Disconnecting the network attachment cable from the client
- Powering off the client
- Shutting down the client
- Closing or ending the open windows on the client

For 5250 workstations, additional options include:

- Powering off the display
- Disconnecting the cable from the display
- Having the user issue a cancel or end request to terminate the active function

12.5.1.3 Server Failures

Failures that can be forced on the server side (that is, the AS/400 side), include:

- Ending jobs immediately (ENDJOB *IMMED)
- Ending subsystems immediately (ENDSBS *IMMED)
- Having your function active when an ENDSYS or PWRDWNSYS is done
- Attempting to connect from the client when the necessary server jobs are not active on the AS/400 system
- For system-to-system connections (that is, AS/400-to-AS/400), requesting an immediate shutdown on one of the AS/400 systems to test the recovery on the second

Note: If system-to-system testing applies to your function, consider having a large system (N-way) or multiprocessor system drive a smaller system. This can place large stress on the small system, but is useful in shaking out problems due to a difference in the timing of an increased stress load.

12.5.2 ERP Testing Tips

Focus on the connect, disconnect, and error recovery paths. Experience shows that things work well once the system and network are up and running, but doing the start-up, takedown, or error recovery is where problems tend to surface.

Recommendation

Besure to repeat the tests again. In error recovery testing, you are looking for timing windows that cause some unusual path to be taken. The repetitive re-application of a test often surfaces some problems.

12.5.3 Problem Determination

Finally, you need to know where to look for error messages. The following lists the places to look:

- QSYSMSG message queue (if it exists)
- QSYSOPR message queue
- QHST
- Joblogs
 - QSYSARB
 - QLUS
 - QCMNARBnn jobs
 - QSYSCOMM1

- QPASVRP job
 - Subsystem joblogs
- Problem logs
- Product activity logs
- LIC logs

Refer to *Communications Management*, SC41-5406, and *CL Programming*, SC41-5721, for further information on problem determination and managing messages.

Chapter 13. Network Availability—TCP/IP Considerations

With the generation of e-business, TCP/IP network stability is becoming more and more critical. From a platform perspective, there are network specific configurations that are critical for TCP/IP access to your AS/400 systems. We recommend that you periodically save these configurations as part of your normal AS/400 backup procedures.

13.1 Saving TCP/IP Configurations

TCP configuration objects stored in multiple locations partly depend on how they are used (by the base operating system or by TCP applications). Files used for basic TCP functions, as well as the IP stack, are shipped in the QATOC* files and QUSRSYS libraries. These objects are shipped to all customers as part of the base operating system product 57nn-SS1.

Objects used by TCP applications reside in the QUSRSYS library with names beginning with QATM*. These files are shipped with the TCP program product 57nn-TC1.

Note: When TCP is installed, the QATM* files initially reside in the QTCP library. Therefore, you do not need to back up the QTCP library to preserve an AS/400 TCP customized environment. You also do not need to restore the QTCP library to regain configuration. The production files reside in the QUSRSYS library.

We recommend that you save and restore all TCP/IP configuration files as a group. Most files are in the QUSRSYS library. However, beginning in V4R1, the TCP product stores files in the integrated file system as well. Some products, such as domain name services, *only* store their files in the integrated file system structure.

Note: Refer to Chapter 15, “Backup and Restore for Integrated File System Objects” on page 221, for information about backing up and restoring objects in the integrated file system.

TCP functions use many logical files, which are defined on physical files. The names of the logical files used in conjunction with TCP are:

- QATMSMTLA—SMTP logical file for system alias table
- QATOCLHST1—Host table logical view 1
- QATOCLHST2—Host table logical view 2
- QATOCLIFC—Interface logical file
- QATOCLN1—Network table logical view 1
- QATOCLN2—Network table logical view 2
- QATOCLPORT—Port logical file
- QATOCLP1—Protocol table logical view 1
- QATOCLRT2—Route table logical view 2
- QATOCLS1—Services table logical view 1
- QATOCLS2—Services table logical view 2
- QATOCMODL1—Modem information table logical view 1

All of these files reside in the QUSRSYS library.

Saving only the physical files and restoring them can cause problems. That is because different functions use the logical files to access the data stored in the

physical files. When only a physical file is restored, the logical files point to a renamed physical file. The restore database functions create the physical file to maintain the indexes to the logical files.

You can benefit from saving and restoring configuration files as a group. For example, saving and restoring only a subset of the files can cause problems when activating TCP/IP for processing. Saving and restoring as a group allows you to maintain any dependencies between the files.

To save all TCP/IP files as a group, enter the following commands:

```
SAVOBJ OBJ(QATOC*) LIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*FILE)
SAVOBJ OBJ(QATM*) LIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*FILE)
```

To restore all TCP/IP files as a group, enter the following commands:

```
RSTOBJ OBJ(QATOC*) SAVLIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*FILE)
RSTOBJ OBJ(QATM*) SAVLIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*FILE)
```

These commands ensure that the TCP files shipped with the operating system (in the QATOC* files and QUSRSYS libraries) and with the program product (57nn-TC1) are included.

For more information about the file structure of the TCP product, see Appendix B of the *TCP/IP Configuration and Reference*, SC41-5420-01.

13.2 Saving Integrated File System Objects

In addition to the files in the QUSRSYS library, objects related to TCP components are created in the integrated file system. They are placed under the integrated file system /QIBM/UserData/*.* directory.

There are four specific directories relative to the configuration and data files in the TCP/IP integrated file system:

- For the Dynamic Host Configuration Protocol (DHCP), save the directory:
/QIBM/UserData/OS400/DHCP/*.*
- For domain name services (DNS), save the configuration files in the directory:
/QIBM/UserData/OS400/DNS
- To save the configuration graphic files created by using the DNS user interface, save the directory:
/QIBM/ProdData/OS400/DNS

ProdData contains DNS related objects used during the configuration process. ProdData should not (but can) have user data stored in it. If this directory is destroyed, re-installing the DNS programs recreates the necessary configurations.

Note: Since the file names for the domain name services are variable (they depend on the domain name), it may be easier to specify files that do not need to be backed up, such as:

- RUNDEBUG
- STATISTICS
- PID
- QUERYLOG
- Anything in the TMP subdirectory

The DHCP files also store their configuration and data files in the integrated file system file structure in the directory:

/QIBM/UserData/OS400/DHCP

Therefore, the save and restore commands you use to back up this environment can refer to these files *generically* to ensure a proper backup and restore process. By the term *generically*, use:

/QIBM/UserData/OS400/DHCP/*.*

For the Web product, specifically Internet connection servers, objects are stored both in the QUSRSYS library and in the integrated file system directory that stores digital certificates:

/QIBM/UserData/WWW

For the Post Office Protocol (POP) mail server, mail is kept in the directory:

/QTCPTMM/MAIL

Mail is normally in this directory until users log on and receive it. For recovery purposes, back up this directory so that the entire mail environment is restored. This operation also benefits the casual user who keeps mail items on the server after reading it and for those users that are away for extended periods of time.

Note: It is worth repeating that, just as with the Domain Mail Services, future TCP (and other) products will utilize the integrated file system structure for its objects. Take this important consideration into account when devising your backup and recovery strategy.

13.3 TCP/IP Tips

This section describes some tools that can assist the system administrator in managing and trouble shooting the TCP/IP environment.

Due to the heritage of TCP/IP and the AS/400 system, many AS/400 TCP/IP users do not realize that AS/400 CL commands for TCP/IP functions may have equivalent TCP industry standard AS/400 alias CL commands. For example:

TCP Alias Command	AS/400 Native Command	Description of Command
FTP	STRTCPFTP	Start TCP/IP File Transfer
LPR	SNDCPSPLF	Send TCP/IP Spooled File
NSLOOKUP	STRDNSQRY	Start DNS Query
PING	VFYTCPCNN	Verify TCP/IP Connection

Table 17 (Page 2 of 2). Command Aliases

TCP Alias Command	AS/400 Native Command	Description of Command
TELNET	STRTCPTELN	Start TCP/IP Telnet

Your AS/400 system is equipped with a number of tools to help identify and resolve problems in the TCP configuration or network.

The OS/400 NETSTAT command allows you to see and control the interfaces, routes, and IP connections that are active on the system. This section outlines other common tools.

13.3.1.1 PING

The PING command is the most commonly used diagnostic tool in a TCP environment. PING stands for Packet Internet Groper. It verifies that connections are possible to other machines. You can run PING on a system to a name if you have an entry in your host table or DNS server for it. Otherwise, you must perform PING on the IP address. If you run PING on the IP address, use apostrophes around the address. For example,

```
PING '051.005.026.001'
```

13.3.1.2 LOOPBACK

The LOOPBACK command is a common TCP/IP function that allows you to test your TCP/IP applications (and most of the interfaces) by accessing the server functions on your own machine. Nothing is transmitted on the network using this command.

In other words, you can run TELNET on your own machine with the command string:

```
TELNET LOOPBACK
```

A 5250 style sign-on display appears shortly after issuing this command.

Note: Use the ATTN key to bring up the Send TELNET Control Functions menu, and end your Telnet session to return to your original session.

You can also run FTP and LPR on your own system. Be aware that you must send a file to a different library or folder than where it resides. For LPR, you must send the print file to a different output queue.

Using the IP address or the name of the system, you can start a session with your own machine. If you are on a Token-Ring network, the packets are sent onto the network and returned on your system. If you can perform PING on your own IP address on the Token-Ring LAN, the AS/400 system is not the source of the problem.

Ethernet can use full-duplex mediums with Ethernet as a switched hub. However, since Ethernet is not a true full-duplex medium, it does not send and receive at the same time. While PING sends packets out on the network and you access your own machine on an Ethernet network, the packets are actually received through an internal software interface in the Ethernet IOP code. Even if you successfully PING your own IP address, there is a possibility that the AS/400 system is the failing component.

For example, if you PING your own address and no packets arrive at the network hub, this almost always indicates a problem in the network between the AS/400 system and the hub. It may indicate a cable problem, or that the port on the hub is disabled.

13.3.2 Common TCP/IP Problems

The following table outlines common symptoms or problems in a TCP environment and where they are typically found. The Things to Check column assumes that:

1. Libraries for the TCP 57nn-TC1 product are installed.
2. QTCP library is in the library list.

<i>Table 18 (Page 1 of 2). TCP Problem Checklist</i>	
Symptoms	Things to Check
TCP/IP does not start.	<ul style="list-style-type: none"> • Line problems exist. • There is a failing Ethernet interface (likely a cable, transceiver, or hub problem). • Check the QTCPIP job log for errors.
You cannot configure TCP/IP.	<ul style="list-style-type: none"> • The user profile doing the configuration does not have *IOSYSCFG special authority.
You can run PING to your own address but not to an address on another machine.	<ul style="list-style-type: none"> • There is an incorrect or missing routing entry. • The subnet mask value is incorrect. • A network problem is indicated.
Other machines can run PING to your system, but they cannot run TELNET to your system.	<ul style="list-style-type: none"> • The QAUTOVRT system value is set to 0 or is not a large enough value. • The QPADEV* device is disabled or hung. • The Telnet server is not active. • The QTGTENETS server job is not started. • The <i>from</i> or <i>to</i> system is a DBCS system.
When TELNET is used to connect to other systems, the user's screen shows distorted information.	<ul style="list-style-type: none"> • The maximum number of screen attributes is exceeded in VT mode • You need to use <i>Display Attributes = *NO</i> on the TELNET command. • You need to use ASCII attributes, such as *AUTOWRAP, on the TELNET command. • The VT machine is sending unsupported special attributes. • You may be missing a PTF to force the AS/400 system into VT220 mode.

Table 18 (Page 2 of 2). TCP Problem Checklist

Symptoms	Things to Check
File transfers do not complete.	<ul style="list-style-type: none"> • The maximum file size supported for FTP on the operating system is exceeded. • The short and long names in the host table or DNS server are not present, causing slow performance. • There is insufficient storage on the source or target system. • The client machine needs to use the TIME command to set the timeout value. • There is a possible software defect. • There is a CCSID or multinational translation involved, which may indicate normal operations.
Records are truncated together using file transfer.	<ul style="list-style-type: none"> • Use BLOCK or BINARY to send data to other EBCDIC machines. • This is a normal condition when the AS/400 system sends text files. FTP truncates trailing blanks and inserts a CR/LF character at the end of each record. You can avoid this by making the last character of each record non-blank. Or try using the BLOCK mode for transmitting data to other EBCDIC machines.
LPR fails because of an unresolved address.	<ul style="list-style-type: none"> • The domain name is not configured.
You can LPR to your own system but not to a remote system or printer.	<ul style="list-style-type: none"> • The problem is at the remote system or printer. • Your AS/400 system may not be authorized by the target system to send print files. Most UNIX machines require that an LPR user is either a <i>trust host</i> or that it is listed in a special file authorizing LPR operations.

For more information on problem determination, see the Trouble Shooting Appendix in the *TCP/IP Configuration and Reference*, SC41-5420-01.

Chapter 14. Availability Options with Hypertext Transfer Protocol

With the growing development of e-business strategies, AS/400 support demands that you understand Hypertext Transfer Protocol (HTTP) server configuration and HTTP implementation around e-business solutions. This chapter addresses which AS/400 HTTP components to back up to ensure complete restoration if needed. We assume that you are familiar with Web-serving terminology prior to reading this chapter.

14.1 Backing Up HTTP Files

We recommend that you back up the following files periodically as part of your normal AS/400 system backup procedures:

1. Configuration and attribute information
 - QUSRSYS/QATMHTTP
 - QUSRSYS/QATMHTTTPC
 - QUSRSYS/QATMHTTPA
 - QUSRSYS/QATMHINSTC
 - QUSRSYS/QATMHINSTA
2. If access control is used, include the following items associated with HTTP server protection setups:
 - Validation list objects associated with HTTP server protection
 - Group files
 - ACL files
3. Web content
4. If using a secure connection, include the backup of:
 - Key-ring files (the name of the file where you want to store public-private key pairs that the server can use for secure information)
 - Signed certificates and, optionally, certificate requests

14.2 Where to Find HTTP Server Protection Setup

When running a Web server on a multi-user system, such as the AS/400 system, protect the data that you serve with the HTTP server properly on the system as you would with any critical data. The AS/400 objects that affect the security of the HTTP server are:

- QTMHHTTP—Default user profile for server
- QTMHHTTP1—Default user profile for CGI programs such as Net.Data
- QUSRSYS/QATMHTTTPC—HTTP configuration file
- QUSRSYS/QATMHTTP—HTTP server attribute file
- Object authority for AS/400 objects to be accessed by the server

A decision must be made about how to set up object authority for HTML documents served by the HTTP server. The QTMHHTTP server profile needs read and execute access to these documents. To achieve a good balance of performance and ease of maintenance, create directories with the default authorities as shown in Figure 64 on page 218.

```

                                Create Directory (CRTDIR)

Type choices, press Enter.

Directory . . . . . harm/webdocs/public

Public authority for data . . . *RX      Name, *INDIR, *RWX, *RW
Public authority for object . . *NONE   *INDIR, *NONE, *ALL...
      + for more values
Auditing value for objects . . . *SYSVAL *SYSVAL, *NONE, *USRPRF

```

Figure 64. Object Authority for HTML Directories

Creating the directories automatically sets authority for new objects added to the directory. Assuming the Webmaster user creates new stream files in the *harm/webdocs/public* directory, the object authority for the *pubhom.htm* document is set as shown in Figure 65.

```

Object . . . . . : /harm/webdocs/public/pubhom.htm
Owner . . . . . : WEBMASTER
Primary group . . . . . : *NONE
Authorization list . . . . . : *NONE

Type options, press Enter.
  1=Add user  2=Change user authority  4=Remove user

      Data      --Object Authorities--
Opt  User      Authority Exist Mgt Alter Ref
-----
 *PUBLIC  *RX
WEBMASTER *RWX           X   X   X   X

```

Figure 65. Object Authority for Documents to be Accessed by the Read-Only Server

Back up the directories in the integrated file system to ensure the complete recovery of the HTTP configuration on your AS/400 system.

For additional information and examples on this topic, see the section on User Authority Setup in the redbook *AS/400 E-Commerce Internet Connection Servers*, SG24-2150.

14.3 HTTP Access Control and Management

The system administrator should become familiar with two log files to manage an HTTP environment:

- Access log—Contains the host name and IP address of the requester, the time of access, the request, the number of bytes transferred, and the return code
- Error log—Contains the error messages, short name and IP address of the requester, the request, and time of error

To limit the size of the access log, do not log requests from selected IP addresses.

There are two formats for these log files:

- The common format—This is the standard format. It can be stored in the integrated file system or in the QUSRSYS library.
- The data description specifications (DDS) format—This format is stored in a physical file in the QUSRSYS library.

For a detailed layout of these log files, refer to the *Internet Connection Server and AS/400 Webmaster's Guide*, GC41-5434.

14.4 HTTP Problem Determination

Normal problem diagnostic methods apply for HTTP problems. To determine the problem:

- Check the configuration using the WRKHTTPCFG command. Look for highlighted lines indicating errors.
- Issue the STRTCPSVR SERVER(*HTTP) HTTPSVR(instance-name) command. Check the job log.
- Check the QEZJOBLOG output queue for server jobs' job logs.

14.4.1 HTTP Log File Setup Using the DDS Format

In this section, we describe the steps to set up your server instance to log access and errors to a file. We chose the DDS format for this scenario. To enable logging, set:

- The global log file configuration settings
- The access log file configuration settings
- The error log file configuration settings

To set up the global log, follow these steps:

1. From the Configuration and Administration Forms page, click on Global Log File Configuration Settings. This form allows you to mark the settings that apply to both the access and error log files.
2. Select DDS for the Log file format prompt. This puts the log files in an AS/400 physical file, which are accessed using normal AS/400 query tools.
3. Click Apply.
4. On the next page, click Configuration Page.

Then, set up the access log.

1. From the Configuration and Administration Forms page, click on Access Log File Configuration. This form allows you to set up the access log file.
2. Enter a name that is meaningful for Access Log Path and Name, such as ACCLOG. This creates a file in the QUSRSYS library for the access log. If you want to use a log file in the integrated file system, enter the complete path here (for example, /home/logs/access), and choose the Common Log Format option instead of DDS.

Note: For integrated file system placement, use the Common format rather than DDS.

3. Click Apply. We will not exclude any requests from the access log in this scenario.

4. On the next page, click Configuration Page.

Next, set up the error log.

1. From the Configuration and Administration Forms page, click on Error Log File Configuration. This form allows you to set the location of the error log file.
2. Enter a name for the Error Log Path and Name, such as:
ERRLOG
3. Click Apply.
4. On the next page, click Configuration Page.

Now, we are back to the Configuration and Administration Forms page. Your settings for the log files are entered. When the instance is running, a member is created in each log file every day using a timestamp, such as "access.Q.0980526" for May 26, 1998. The member contains the access information and errors for that day.

For the Access Log File Configuration panel, you can specify the IP addresses or host names to not log. For example, to not log accesses from the IP addresses of *100.27.22.** or **.ibm.com*, specify:

*100.27.22.**

or

**.ibm.com*

14.5 For More Information

For more information about HTTP and the AS/400 system, refer to the following URL address:

<http://www.as400.ibm.com/tstudio/workshop/webbuild.htm>

Chapter 15. Backup and Restore for Integrated File System Objects

The addition of the integrated file system structure in V3R1 provides a flexible structure for managing and retrieving information. Objects of all types can be referenced through the integrated file system structure.

This chapter discusses save and restore commands using the integrated file system interface. The integrated file system is becoming an increasingly important area for the system administrator to understand and manage. This is especially true as licensed program products and the operating system more frequently store data in file systems other than libraries and folders.

We begin with a discussion of the integrated file system and, then, advance to the various licensed program products that use this interface, such as:

- User-defined file system
- Document library objects
- Domino for AS/400
- Windows NT
- Lotus Notes on the Integrated PC Server
- NetWare on the Integrated PC Server
- OS/2 Warp Server for AS/400
- PC client
- Firewall

15.1 The Integrated File System

The integrated file system integrates existing AS/400 library objects, folders, documents, and shared folders with new object types, such as STMF and directories, into a single hierarchical file system. Traditional AS/400 file systems, such as libraries (QSYS.LIB) and folders (QDLS), are integrated with new file systems in a single hierarchical file system. The integrated file system is part of the AS/400 operating system (OS/400), and supports stream I/O and storage management. This hierarchical file system is similar to personal computers and UNIX operating systems, yet provides for an integration of traditional AS/400 file structures.

The following diagram shows the integrated file system structure.

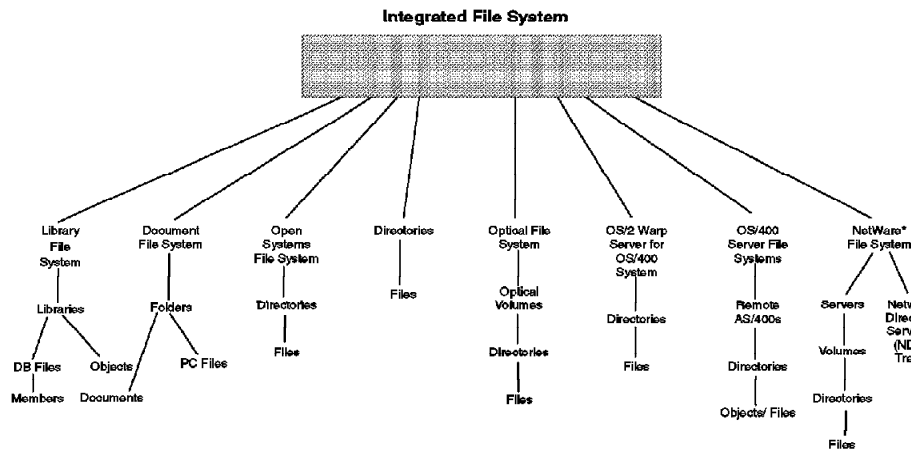


Figure 66. Integrated File System Structure

The integrated file system consists of the following file systems:

- QSYS.LIB (the library file system)

The QSYS.LIB file system supports the AS/400 library structure. It provides access to database files and all AS/400 object types that are managed by library support. It supports programming languages and facilities that operate on the database files. This file system has extensive administrative support. It also supports stream I/O to user spaces, source physical file members, and database members that are program defined.

- QDLS (the document library services file system)

The QDLS file system supports the folder structure, and provides access to documents and folders. This system is integrated with OfficeVision/400, and stores data in stream files.

- QLanSrv (the OS/2 Warp Server for AS/400 file system)

Note: The LAN Server/400 file system is now called OS/2 Warp Server for AS/400. Its directory is still called QLanSrv.

The QLanSrv file system provides access to the same directories and files that are accessed through the LAN Server/400 licensed program. It allows users of the OS/400 file server and AS/400 applications to use the same data as LAN Server clients. The LAN Server/400 file system is available only when the LAN Server licensed program is installed on the system. The data is stored in stream files.

- QOpenSys (the open systems file system)

The QOpenSys file system is compatible with UNIX-based open system standards such as POSIX and XPG. An important distinguishing feature of this system from the root-file system is that it supports case-sensitive names. It also supports hard links and symbolic links and has good performance for

stream files. It supports long names and NLS enabled names. The names are stored internally in ISO 10646 (UCSC Level 1) or Unicode format.

- Root (the “/” file system)

The root file system takes full advantage of the stream file support and hierarchical directory structure of the integrated file system. It has characteristics of both DOS and OS/2 file systems. It has similar characteristics to the QOpenSys file system.

- QNetWare (the NetWare file system)

The QNetWare file system provides access to local and remote data and objects that are stored on a server that runs Novell NetWare 3.12 or 4.10. It also offers access to NetWare Directory Services (NDS) objects. The QNetWare file system provides dynamic mounting of NetWare file systems into the local name space. As opposed to the QOpenSys file system, QNetWare is not case sensitive.

- UDFS (the user-defined file system)

A UDFS file system resides in the ASP of the user’s choice. The user creates and manages these file systems. They can even choose to specify if the attributes are case-sensitive when creating the UDFS. The user must mount the UDFS into the local name space for visibility and accessibility. A user-defined file system is represented as a block special file (*BLKSF). It supports a graphical interface, and once mounted, provides the same support as the root and QOpenSys file systems. Refer to Section 15.3, “User-Defined File System” on page 231, for an example.

Note: The Add Mounted File System (ADDMFS) and MOUNT commands make the objects in this file system accessible to the integrated file system name space. To mount a UDFS, specify *UDFS for the TYPE parameter on the ADDMFS command.

- QOPT (the optical file system)

The QOPT file system provides access to direct attached optical libraries and CD-ROM. It stores data in stream files.

- QFileSvr.400 (the file server file system)

The QFileSvr.400 file system allows transparent access to file systems that reside on remote AS/400 systems. Think of it as a client that acts on behalf of users to perform a file request. QFileSvr.400 interacts with the OS/400 file server on the target system to perform the actual file operation. For a first-level directory (which actually represents the root (/) directory of the target system), the QFileSvr.400 file system preserves the same uppercase and lowercase attributes when QFileSvr.400 searches for names. For other directories, case-sensitivity depends on the specific file system that is accessed.

15.2 Save and Restore for the Integrated File System

The release of V3R1 introduced the SAV and RST commands to save and restore data stored in the integrated file system. These commands are available when the SAVE and RESTORE menu options are selected. These menu options, including the SAV command, allow you to choose the following options from the SAVE menu:

- Option 21, which saves the entire system

- Option 23, which saves all user data

Similarly, the corresponding RESTORE menu offers the following options, including the RST command:

- Option 21, which restores the entire system
- Option 23, which restores all the user data

15.2.1 When to Use the SAV Command

The traditional save and restore operating system commands remain the preferred way to back up and recover objects residing in the QSYS.LIB and QDLS file systems on the system.

Note: The intent of the SAV and RST commands are *not* to replace all existing commands where they (the commands) are available. Instead, users of the integrated file system can use the SAV and RST commands to save data in the integrated file system.

Recommendation

Beginning with V3R1, you must *use the SAV command* to save data outside QSYS.LIB and QDLS, or you cannot completely back up the system.

Using the SAV and RST commands, you can save and restore:

- Client Access/400 objects in directories
- LAN Server/400 objects in directories
- NetWare objects in directories
- Domino Notes for AS/400
- Other licensed program objects in directories
- Customer and vendor applications in directories

With the AS/400 system, you can perform save operations on objects that are changed either from a specific date and time or since the last save operation. Use the CHGPERIOD parameter on the SAV command to perform this operation.

Refer to “Is Your Entire System Backed Up?” in the April 1998 issue of NEWS/400 for more information on saving the integrated file system with the SAV and RST commands. Information about NEWS/400 and past issues can be found at:

<http://www.news400.com/>

Note: The traditional save and restore operating system commands will remain the preferred way to back up and recover objects on the system. The intent of the SAV and RST commands are *not* to replace all existing commands. Instead, the intent is to allow new users of the integrated file system to use the new commands to be able to save data in the integrated file system.

15.2.2 Considerations When Saving Across Multiple File Systems

Because different file systems support different object types and naming conventions, you cannot save by object name or object type. You can save all objects from all file systems or omit some file systems.

The following examples demonstrate each approach.

- For saving all objects in all file systems, except QSYS.LIB and QDLS, enter:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/*')
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) UPDHST(*YES)
```

Note: This is the recommended approach.

- For saving all objects on the system, enter:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/*')
```

Attention

Although this is a valid combination, do *not* use the SAV command to save the entire system. See Section 15.2.3.1, “Tips for Performance and Reducing the Time to Save” on page 227, for more information.

- For saving all objects in the file system, except the QSYS, QDLS, and one or more other file systems, enter:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/*')
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)
('/other-values' *OMIT)) UPDHST(*YES)
```

Note: When you indicate SAV OBJ('/*'):

- The system must be in a restricted state.
- *SAVSYS or *ALLOBJ special authority is required.
- DEV cannot be a save file or diskette file.
- You must have VOL(*MOUNTED).
- You must specify SEQNBR(*END).

Values for other parameters of the SAV command are supported only for some file systems. Choose the lowest common denominator. Refer to *Backup and Recovery*, SC41-5304, for examples.

Specify the following values for the other parameters:

DEV	Tape device only (no save file nor diskette file)
SUBTREE	*ALL
SAVACT	*NO
OUTPUT	*NONE
LABEL	*GEN
USEOPTBLK	*NO
SYSTEM	*LCL
CHGPERIOD	default
PRECHK	*NO
UPDHST	*YES

When saving objects from the integrated file system, objects can be locked that either prevent objects from being saved or slow the save down. To avoid locks, perform saves of integrated file system objects when users are not on the system and the system is quiesced.

Should a lock occur, it is not readily apparent who is holding the lock, as the Work With Object Lock (WRKOBJLCK) command does not support integrated file system objects. You may be able to determine what jobs or objects are involved in an integrated file system object lock by deciphering the output from an API.

Enter the following commands and examine the resulting spooled file for names that may lead you to a job as the holder of the lock. Any action taken to free up the lock should be done with caution. If you need assistance in either breaking

the lock or interpreting the output from the API, contact IBM Support Line for additional services. The commands are:

```
CALL QPOFTPOS *DUMP
CALL QPOFTPOS 'XXXXXX'
```

The value XXXXXX is a term found in the QPOFTPOS dump.

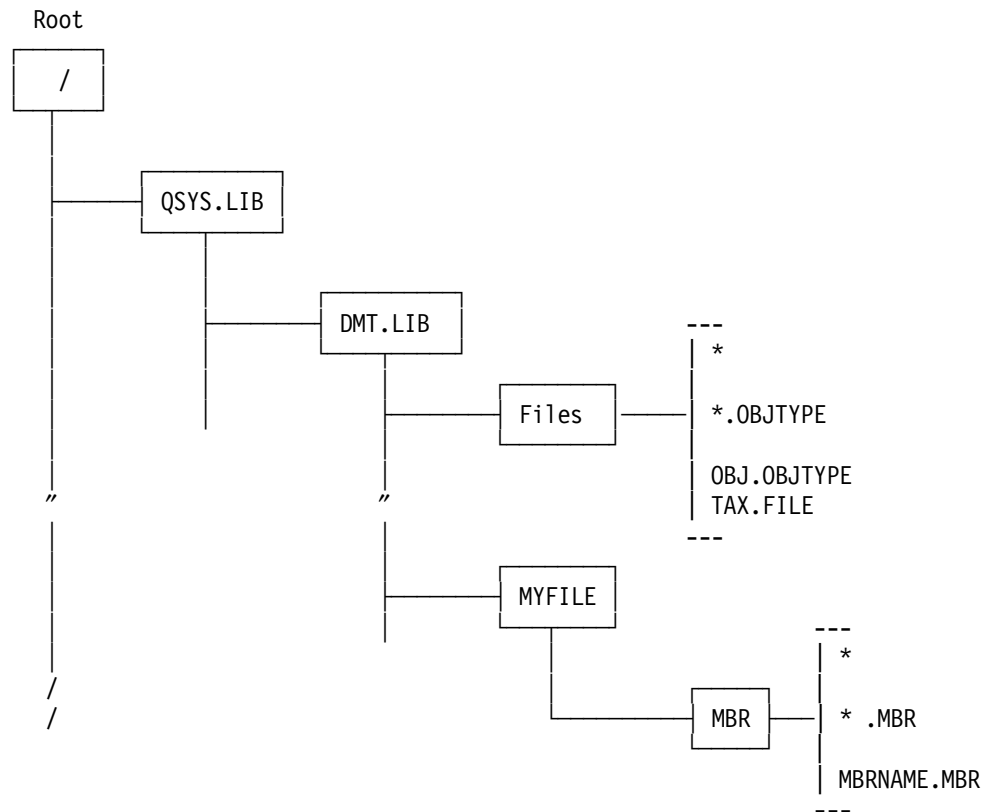
15.2.3 Considerations when Saving Objects from the QSYS.LIB File System

The path name element for QSYS.LIB objects must be in the *name.objecttype* format, such as:

```
'/QSYS.LIB/DMT.LIB/TAX.FILE'
```

The object name and object type are separated by a period (.). Objects in a library can have the same name if they are different object types. Therefore, specify each object type to uniquely identify the object.

The following figure demonstrates a graphical method of looking at the structure.



Structure of QSYS.LIB

Figure 67. The Structure of the QSYS.LIB

When saving objects from the QSYS.LIB file system, these restrictions apply:

- The *OBJ* parameter must have only one name.
- The SAV command does not have the ACCPTH (save access path) keyword, which means by default, these commands do not save access paths.

Therefore, the SAV command does *not* save any access paths for the QSYS.LIB file system (including all system and user libraries).

- To save a library, enter:

```
SAV    DEV('/ QSYS.LIB/tape-device-name.DEVD')
        OBJ(' QSYS.LIB/library-name.LIB')
```

- To save all objects in a library, enter:

```
SAV    DEV('/ QSYS.LIB/tape-device-name.DEVD')
        OBJ(' QSYS.LIB/library-name.LIB/*')
```

15.2.3.1 Tips for Performance and Reducing the Time to Save

For QSYS objects, the SAV command maps the save request to one of the existing save commands (SAVLIB, SAVOBJ, or SAVCHGOBJ and so forth). This effectively puts another layer between the SAV and the equivalent SAVxxx command, which can result in performance degradation.

SAV OBJ('/**') is not the recommended way to save the entire system. Use AS/400 save commands, such as option 21 from the SAVE Menu, or similar commands to save the entire system. Refer to Chapter 5, "Save and Restore for Availability and Recovery" on page 49, for information on using the SAVE menu and the SAV/RST commands.

You get the best SAV performance by using the 3590 or 3570 tape devices. This is also true for other commands, not just the SAV command. It is beneficial if the tape drive is not on the same IOP as the load source. In other words, put the 3590 on a separate bus. Consider this option especially if you have a large amount of integrated file system data to save.

Recommendation

If you have a large number of stream files (about 18GB or more), you may find that the performance of the SAV command is considerably slower. A rule of thumb is that if you save a lot of small objects (as compared to a few larger objects), your performance can suffer.

Note: At the time of writing this redbook, there is a PTF to reduce an unnecessary VLIC log generation that impedes the cleanup process and delays the save. The PTF for V4R1 is MF17584, and the PTF for V4R2 is MF17910. The Authorized Problem Analysis Request (APAR) number relating to this PTF is MA17273. Contact your IBM support representative for more information and to confirm that the PTF named is the latest recommendation.

The QSYS.LIB file system supports "externally described" database files. Externally described refers to data in fields and records of a file that are described outside of the program that processes the file. Such files include those created by DDS, IDDU, or DB2 for the AS/400 licensed program. You can use the SAV command to save these externally described files. You can perform SAV on any object on which you can perform a SAVOBJ.

Attention

Currently users gain access to the QSYS.LIB file system by using the integrated file system interface. To restrict their access to QSYS.LIB, use the QPWFSESERVER authorization list.

To restrict one user, change the authority specified in the QPWFSESERVER authorization list to *EXCLUDE. To restrict all users, change *PUBLIC authority to *EXCLUDE. To grant access to a specific user, change the authority to the QPWFSESERVER authorization list to *USE by using the GRTOBJAUT command, or *RX by using the CHGAUT command.

If the authority of QPWFSESERVER type *AUTL changes while the SERVER subsystem is active, end and restart it to ensure that Client Access/400 users detect the change.

If you are on V3R1, make sure to apply PTF SF24822 or a superseding PTF to restrict QSYS.LIB access.

15.2.4 Considerations when Restoring across Multiple File Systems

Because different file systems support different object types and naming conventions, you cannot specify a restore by object names and object type. However, you can restore all objects from all file systems or omit file systems.

Use these valid combinations:

- Restoring all objects on the system

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/*')
```

Note: This is *not* the recommended method to restore, just as SAV OBJ('/*') is not the recommended way to save. Use the RST command with the OMIT parameter. For example:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))
```

- Restoring all objects, except QSYS.LIB and QDLS

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/*') ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))
```

- Restoring all objects, except QSYS.LIB, QDLS, and one or more other file systems

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/*')
    ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)
    ('XYZ.LIB' *OMIT))
```

Note: When you specify RST OBJ('/*'):

- The system restores only objects saved by SAV OBJ('/*').
- The system must be in a restricted state.
- You must have *SAVSYS or *ALLOBJ special authority.
- You cannot specify diskettes, savf, or optical media for DEV.
- You must specify NBR(*SEARCH).

15.2.5 Considerations when Restoring Objects from the QSYS.LIB File System

When you use the RST command to restore objects to the QSYS.LIB file system, these restrictions apply:

- Give the OBJ parameter only one name.
- Specify objects that are allowed on the RSTOBJ command. For example, the RST of user profiles is not allowed because it is not allowed on the RSTOBJ command.
- For QSYS.LIB file system libraries, objects QDOC and QSYS cannot be restored with RSTLIB because QSYS contains system objects and QDOC contains documents. Therefore, you cannot use the RST command to restore the following documents and objects:

```
QDOC  QDOCnnnn  QRECOVER  QRPLOBJ
QSRV  QSPL      QSYS       QTEMP
```

The following table compares the RST and equivalent RSTxxx commands.

<i>Table 19. RST Command Versus the RSTxxx Commands</i>	
Object Parameter on RST Command	Equivalent RSTxxx Command
OBJ('/QSYS.LIB/LIBA.LIB')	RSTLIB SAVLIB(LIBA)
OBJ('/QSYS.LIB/LIBA.LIB/*')	RSTOBJ SAVLIB(LIBA) OBJ(*ALL) OBJTYPE(*ALL)
OBJ('/QSYS.LIB/LIBA.LIB/*.object-type')	RSTOBJ SAVLIB(LIBA) OBJ(*ALL) OBJTYPE(object-type)
OBJ('/QSYS.LIB/LIBA.LIB/object-name.object-type')	RSTOBJ SAVLIB(LIBA) OBJ(object-name) OBJTYPE(object-type)
OBJ('/QSYS.LIB/LIBA.LIB/file-name.FILE/*')	RSTOBJ SAVLIB(LIBA) OBJ(file-name) OBJTYPE(*FILE)
OBJ('/QSYS.LIB/LIBA.LIB/file-name.FILE/*.MBR')	RSTOBJ SAVLIB(LIBA) OBJ(file-name) OBJTYPE(*FILE)
OBJ('/QSYS.LIB/library-name.LIB/file-name.FILE/*MBR')	RSTOBJ SAVLIB(LIBA) OBJ(file-name) OBJTYPE(*FILE) FILEMBR((*ALL (member-name)))

Use the RST command's new-name parameter element to rename an object in a directory, restore an object to a different directory, or restore an object to a different library. Consider these examples:

- File Operations:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
  OBJ('/DBSDIR/FILEB' *INCLUDE '/DBSDIR/FILEX')
```

In this example, FILEX is created in the DBSDIR directory. The data that was saved with FILEB is restored to FILEX. If FILEB still exists on the system, it remains unchanged.

- Directory Operations:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
  OBJ('/DBSDIR/FILE*' *INCLUDE '/LMSDIR')
```

In this example, running this command restores all objects with a name beginning with FILE from the DBSDIR to the /LMSDIR directory.

- Library Operation:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
  OBJ('/QSYS.LIB/LIBA.LIB' *INCLUDE
    '/QSYS.LIB/LIBB.LIB')
```

This library operation restores library A (LIBA) to library B (LIBB). Note that, in this example, Library B does not exist prior to the RST operation. For library operations, the new library is created.

- Restoring all objects:

```
RST DEV('/QSYS.LIB/LIBA.LIB/TESTB.file')
  OBJ('/QSYS.LIB/LIBA.LIB/*' *INCLUDE
    '/QSYS.LIB/LIBB.LIB')
```

The restoring all objects operation restores all objects in LIBA to LIBB. TESTB, in this case, is a save file (*SAVF).

- Operations on all objects of a specified type (for example, the PGM type):

```
RST DEV('/QSYS.LIB/LIBA.LIB/TESTB.FILE')
  OBJ('/QSYS.LIB/LIBA.LIB/*.PGM'
    *INCLUDE '/QSYS.LIB/LIBB.LIB')
```

This RST operation restores all PGM type objects from LIBA to LIBB.

- For database file members, OPTION(*NEW) restores members for new files only:

```
RST DEV('/QSYS.LIB/hlaust.LIB/TESTB.FILE')
  OBJ('/QSYS.LIB/hlaust.LIB/*.FILE'
    *INCLUDE '/QSYS.LIB/hlaustb.LIB')
  OPTION(*NEW)
```

This RST restores all objects from the hlaust directory to the root directory.

Note: In this case, we used default values for the other parameters. However, note that the other parameters must have these values:

```
SUBTREE(*ALL)
SYSTEM(*LCL)
OUTPUT(*NONE)
ALWOBJDIF(*ALL)
```

You can only rename the library, not the objects in the library. Use the WRKLNK command with option 7, Rename or the Rename Object (RNM) command to rename the library.

15.2.6 Considerations when Restoring Objects to the QDLS File System

When you use the RST command to restore objects to the QDLS file system, these restrictions apply:

- The OBJ parameter must have only one name.
- The OBJ and SUBTREE parameters must have one of the following values:

```
OBJ('/QDLS/path/folder-name') SUBTREE(*ALL)
OBJ('/QDLS/path/document-name') SUBTREE(*OBJ)
```

- Other parameters must have these values:

```
SYSTEM *LCL
OUTPUT *NONE
ALWOBJDIF *ALL or *NONE
OPTION *ALL
```

To restore into an existing folder (a folder called FLDX), enter:


```
RST DEV('/QSYS.LIB/LIBA.LIB/TESTB3.FILE')
OBJ('/QDLS/FLDX' *INCLUDE)
ALWOBJDIF(*ALL)
```

This operation restores the FLDX folder into an existing folder, also named FLDX, on the system.

When you restore into an existing folder, create a new folder using the object name parameter on the RST command:

```
RST DEV('/QSYS.LIB/LIBA.LIB/TESTB3.FILE')
OBJ('/QDLS/FLDX' *INCLUDE '/QDLS/FLDY'))
ALWOBJDIF(*ALL)
```

This results in the error message:

```
CPF909C System object name KN3Q241287 already exists.
```

To correct this, rename or delete the existing New Object Name folder FLDX. Now, you can restore into the folder specified by the parameter FLDY.

15.3 User-Defined File System

In any user-defined file system (UDFS), as in the root and QOpenSys file systems, you can create directories, stream files, symbolic links, local sockets, and SOM objects. The UDFS includes the following characteristics:

- A path name to a UDFS must be in the form:

```
/dev/asppXX/udfs-name.udfs
```

Failure to use this convention results in the error message:

```
CPFA0A2 - Information passed to this operation not valid.
```

- A UDFS can be created in a user ASP or the system ASP.
- You can specify case-sensitivity. This is done when you run the command:

```
CRTUDFS UDFS('/dev/QASP01/HL.UDFS')
CASE(*MIXED)
```

Specifying CASE(*MIXED) results in case sensitivity. CASE(*MONO) results in no case sensitivity.

- A UDFS exists in two states: mounted and unmounted. Find a discussion of this important concept in the following section, Section 15.3.1, “Mounting User-Defined File System.”
- OS/400 sees UDFS as a block special file object (*BLKSF). As you create UDFSs, the block special files are created automatically. These files are accessible only by using the integrated file system generic commands, APIs, and the QFileSvr.400 interface.
- Once mounted, UDFS provides the same support as the root and QOpenSys file systems.

15.3.1 Mounting User-Defined File System

The word “mounting” implies making a file system visible and accessible through the integrated file system namespace. Prior to mounting, the contents of the UDFS are not visible or accessible. To make it easier to understand this concept, we offer the following displays and outline the steps to complete it.

Using a directory called hltest, issue a WRKLNK command to see the Work with Object Links display:

```
WRKLNK ('/*')
```

```
Work with Object Links

Directory . . . . : /

Type options, press Enter.
 3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
11=Change current directory ...

Opt  Object link      Type  Attribute  Text
---  ---
   dev_dpmon         STMF
   domino           DIR
   etc               DIR
   example-resources DIR
   glms              DIR
   sample.dsk        STMF
   sample.nsf        STMF
 5   hltest          DIR
   home              DIR

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field      F23=More options

More...
```

Figure 68. Work with Object Links Display (Part 1 of 2)

1. The directory called hltest appears.
2. Next, select option 5 to go to the next level, as shown in the following display. Notice that there are two STMF files in the directory /hltest.

```

                                Work with Object Links

Directory . . . . : /hltest

Type options, press Enter.
  3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type   Attribute  Text
—      jdbcTest.example.C >  STMF
—      jdbcTest.example.C >  STMF

                                                                 Bottom

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field          F23=More options

```

Figure 69. Work with Object Links Display (Part 2 of 2)

3. At this stage, the contents of /hltest are visible and accessible through the integrated file system interfaces.
4. The Display Mounted File System Information (DSPMFSINF) command shows information about a mounted file system. The following display shows information about the /hltest mounted file system.

```

                                Display Mounted FS Information

Object . . . . . : /hltest

File system type . . . . . : "root" (/)

Block size . . . . . : 4096
Total blocks . . . . . : 4993024
Blocks free . . . . . : 669368
Object link maximum . . . . . : 32767
Directory link maximum . . . . . : 32767
Pathname component maximum . . . : 510
Path name maximum . . . . . : No maximum
Change owner restricted . . . . . : Yes
No truncation . . . . . : Yes
Case Sensitivity . . . . . : No

Press Enter to continue.

F3=Exit  F12=Cancel
(C) COPYRIGHT IBM CORP. 1980, 1998.

```

Figure 70. Display Mounted FS Information

Note that the file system type is root (/).

- Now, issue the Create User Defined File System (CRTUDFS) command with the parameters:

```
CRTUDFS UDFS('/dev/QASP01/HL.UDFS')
```

Note that the system ASP (QASP01) is used in this example.

- Mount the UDFS on top of /hltest:

```
ADDMFS TYPE(*UDFS) MFS('/dev/Qasp01/hl.udfs')
MNTOVRDIR('/hltest')
```

- The WRKLNK command output shows the mounted structure with the path of the mounted file system /dev/Qasp01/hl.udfs, which is mounted over the original path of /hltest.

Next, choose option 5 for the next level.

```

Work with Object Links

Directory . . . . : /

Type options, press Enter.
 3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
11=Change current directory ...

Opt  Object link      Type  Attribute  Text
--  -
dev_dpmon      STMF
domino         DIR
etc            DIR
example-resources  DIR
gllms         DIR
dmt.dsk       STMF
dmt.nsf       STMF
5  hltest           DIR          test for hlaust user
home          DIR

More...

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field          F23=More options

```

Figure 71. Work with Object Links Display

- Now it has the contents of /dev/qasp01/hl.udfs.

Note: The object hltest.example.Cre was created some time after step 6. Our example is for an empty directory. Hltest.example.Cre is shown for illustration only.

```

Work with Object Links

Directory . . . . : /hltest

Type options, press Enter.
 3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
11=Change current directory ...

Opt  Object link      Type  Attribute  Text
--  -
hltest.example.Cre > STMF

```

Figure 72. Work with Object Links Display

This display also indicates that the former contents of /hltest directory (such as jdbcTest.example.C as shown in Figure 69 on page 233) is no longer visible and accessible. The original /hltest is “masked” by this UDFS.

9. If you issue the Display Mounted FS Information (DSPMFSINF) command DSPMFSINF OBJ('/hltest'), the following display appears:

```

                                Display Mounted FS Information
Object . . . . . : /hltest
File system type . . . . . : User-defined file system
Block size . . . . . : 4096
Total blocks . . . . . : 4993024
Blocks free . . . . . : 773760
Object link maximum . . . . . : 32767
Directory link maximum . . . . . : 32767
Pathname component maximum . . . : 510
Path name maximum . . . . . : No maximum
Change owner restricted . . . . . : Yes
No truncation . . . . . : Yes
Case Sensitivity . . . . . : No
Path of mounted file system . . . : /dev/Qasp01/h1.udfs
Path mounted over . . . . . : /hltest
Protection . . . . . : Read-write
Setuid execution . . . . . : Not supported
Mount type . . . . . : Not supported
Read buffer size . . . . . : Not supported
Write buffer size . . . . . : Not supported
Timeout . . . . . : Not supported
Retry Attempts . . . . . : Not supported
Retransmission Attempts . . . . . : Not supported
Regular file attribute minimum
time . . . . . : Not supported
Regular file attribute maximum
time . . . . . : Not supported
More...
Press Enter to continue.
F3=Exit F12=Cancel

```

Figure 73. Display Mounted FS Information

Note: Pay attention to the labels “path of mounted file system” and “path mounted over” as shown in the second Display Mounted FS Information display. Note that the file system type changes.

15.3.2 Saving and Restoring an Unmounted UDFS

Follow the rule of thumb to save and restore the user-defined file system. Unmount them with the command:

```
RMVMFS TYPE(*UDFS) MNTOVRDIR('/hltest')
```

As an alternative, use the command:

```
RMVMFS TYPE(*UDFS) MFS('/dev/qasp01/h1.udfs')
```

You cannot specify both MNTOVRDIR and MFS for TYPE(*UDFS).

Note: You can also use the UNMOUNT command with the same parameters.

The Display User-Defined System (DSPUDFS) command enables you to check whether the UDFS file system is mounted. The next display shows an example of the output from this command.

```

                                Display User-Defined FS
User-defined file system . . . : /dev/qasp01/hl.udfs
Owner . . . . . : A97011363
Code page . . . . . : 37
Case sensitivity . . . . . : *MONO
Creation date/time . . . . . : 12/07/97 20:50:01
Change date/time . . . . . : 12/07/97 21:13:07
Path where mounted . . . . . : Not mounted
Description . . . . . : test for user hlaust

                                                                    Bottom
Press Enter to continue.
F3=Exit  F12=Cancel
(C) COPYRIGHT IBM CORP. 1980, 1998.
```

Figure 74. Display User-Defined FS

Note: In this display, case sensitivity is set to *MONO. You may have to set it to *MIXED to enable case sensitivity.

15.3.3 Saving an Unmounted UDFS

When saving a UDFS, save all the objects and the attributes of the UDFS.

To save an unmounted UDFS, enter the command:

```
SAV DEV('/QSYS.LIB/UDFS1.LIB/UDFSSAVF.FILE')
OBJ('/DEV/QASPO1/HL.UDFS') SUBTREE(*ALL)
```

In this example, the UDFS is saved to a previously created save file named UDFSSAVF.

15.3.4 Restrictions when Saving an Unmounted UDFS

There are several restrictions worth noting when saving an unmounted user-defined file system. They are:

- You cannot save individual objects from unmounted UDFSs. You can only save objects from the UDFS all at the same time.
- You cannot view or work with the objects in an unmounted UDFS.
- The default parameter SUBTREE(*ALL) is required. This saves the entire subtree of each directory that matches the object name pattern. The subtree includes all subdirectories and the objects within them.

- The TGTRLS parameter must specify a release value of V3R7M0 or later.

15.3.5 Restoring an Unmounted UDFS

To restore an unmounted UDFS, use the command:

```
RST DEV('/QSYS.LIB/ABCD.LIB/UDFSSAVF.FILE')
    OBJ('/DEV/QASP01/HL.UDFS')
```

If the UDFS does not exist on the system, the *BLKSF is created. If the UDFS exists, objects from the save media overlay objects on the system. If the ASPs from which they were saved do not exist, the UDFSs are not restored from where they were saved.

15.3.6 Restrictions when Restoring an Unmounted UDFS

The following restrictions apply when restoring an unmounted UDFS:

- You cannot view or work with objects in an unmounted UDFS. Therefore, you cannot determine the amount of storage and time required.
- Individual objects cannot be restored to an unmounted UDFS.

15.3.7 Restoring an Individual Object from an Unmounted UDFS

You can restore individual objects from a save media containing unmounted UDFSs. To do so, give a new name to the object being restored. The parent directory of the new name must exist in an accessible file system.

The following steps show an example of how to restore an individual object from an unmounted UDFS.

1. To save the unmounted UDFS, enter: /DEV/QASP01/HL.UDFS with the command:

```
SAV DEV('/QSYS.LIB/HL.LIB/UDFSSAVF3.FILE')
    OBJ('/DEV/QASP01/HL.UDFS')
```

2. To restore the unmounted UDFS to another directory and give it a new name /hltest2/PAYROLL, enter:

```
RST DEV('/QSYS.LIB/hlaust.LIB/UDFSSAVF3.FILE')
    OBJ('/DEV/QASP01/HL.UDFS/hltest.example.Createprocedures.ini'
    *INCLUDE '/hltest2/PAYROLL')
```

15.3.8 Saving a Mounted UDFS

Objects saved from a mounted UDFS are now visible and accessible. The original objects that were contained in the directory before mounting are not saved, but are masked over. The information about the UDFS and ASP also is not saved. In other words, saving objects from a mounted UDFS is like saving a root-file system object.

To better understand the concept of saving a mounted UDFS, follow this example. We use the same directory as in the previous example for mounting UDFS. The /hltest directory, has two objects:

```

Work with Object Links

Directory . . . . : /hltest

Type options, press Enter.
  3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type  Attribute  Text
_      jdbcTest.example.C >  STMF
_      jdbcTest.example.C >  STMF

```

Figure 75. Work with Object Links Display

Mount the specific UDFS: /dev/qasp01/hl.udfs onto /hltest namespace, which appears as shown in the following display:

```

Work with Object Links

Directory . . . . : /hltest

Type options, press Enter.
  3=Copy  4=Remove  5=Next level  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type  Attribute  Text
_      hltest.example.Cre >  STMF

```

Figure 76. Work with Object Links Display

Now, if you save a mounted UDFS and issue a DSPSAVF UDFSSAVF.FILE to view our saved information, the following display appears:

```

Display Saved Objects - Save File

Display level . . . . . : 2
Directory . . . . . : /hltest

Type options, press Enter.
  5=Display objects in subdirectory  8=Display object specific information

Opt  Object                                     Type  Owner      Size
_      hltest.example.Createprocedure >  *STMF  A97011363  8192

```

Figure 77. Display Saved Objects

If you restore the UDFS using:

```

RST DEV(' QSYS.LIB/hlaust.LIB/MOUNT1.file')
OBJ('/hltest' *INCLUDE '/hltest2')

```

and issue the DSPMFSINF command for /hltest2, you lose the attributes. Contrast this against our original /hltest directory as shown earlier.

The following displays illustrate this.


```

                                Display Mounted FS Information

Object . . . . . : /hltest

File system type . . . . . : User-defined file system

Block size . . . . . : 4096
Total blocks . . . . . : 4993024
Blocks free . . . . . : 769659
Object link maximum . . . . . : 32767
Directory link maximum . . . . . : 32767
Pathname component maximum . . . : 510
Path name maximum . . . . . : No maximum
Change owner restricted . . . . . : Yes
No truncation . . . . . : Yes
Case Sensitivity . . . . . : No

Path of mounted file system . . : /dev/qasp01/hl.udfs

Path mounted over . . . . . : /hltest

Protection . . . . . : Read-write
Setuid execution . . . . . : Not supported
Mount type . . . . . : Not supported
Read buffer size . . . . . : Not supported
Write buffer size . . . . . : Not supported
Timeout . . . . . : Not supported
Retry Attempts . . . . . : Not supported
Retransmission Attempts . . . . . : Not supported
Regular file attribute minimum
  time . . . . . : Not supported
Regular file attribute maximum
  time . . . . . : Not supported
Directory attribute minimum
  time . . . . . : Not supported
Directory attribute maximum
  time . . . . . : Not supported
Force refresh of attributes on
  open . . . . . : Not supported
Attribute and name caching . . . : Not supported
Data file code page . . . . . : Not supported
Pathname code page . . . . . : Not supported

```

Figure 78. A Mounted UDFS before Save

```

                                Display Mounted FS Information

Object . . . . . : /hltest2

File system type . . . . . : "root" (/)

Block size . . . . . : 4096
Total blocks . . . . . : 4993024
Blocks free . . . . . : 769659
Object link maximum . . . . . : 32767
Directory link maximum . . . . . : 32767
Pathname component maximum . . . . . : 510
Path name maximum . . . . . : No maximum
Change owner restricted . . . . . : Yes
No truncation . . . . . : Yes
Case Sensitivity . . . . . : No

```

Figure 79. A UDFS Saved while Mounted after Restore

Saving a mounted file system produces this error message in the job log:

CPD3788 No file system information was saved because /hltest is a mounted user-defined file system.

If a UDFS is mounted as read-only when the UDFS is restored, it loses the information shown in Figure 78 on page 239.

15.3.9 Restoring a Mounted UDFS

Objects saved from the mounted UDFSs are restored into the file system of the parent directory to which the objects are restored. The UDFS and ASP information is not restored. Section 15.3.8, "Saving a Mounted UDFS" on page 237 of this chapter explains where the attributes of the restored UDFS and the original UDFS contrast.

Note: In a disaster recovery situation, you must first create the UDFS again, and restore the backup to the newly created UDFS.

15.3.10 Integrated File System Commands

The following outlines a list of integrated file system commands for managing on the UDFS.

<i>Table 20. User-Defined File System CL Commands</i>	
Command	Description
ADDMFS	Adds mounted file system
CRTUDFS	Creates a UDFS
DLTUDFS	Deletes UDFS
DSPMFSINF	Displays mounted file system information
DSPUDFS	Displays UDFS
MOUNT	Mounts a file system
RMVMFS	Removes mounted file system
UNMOUNT	Unmounts a file system

15.4 Document Library Objects

This section describes saving and restoring document library objects (DLOs) using traditional interfaces. You can still use the SAV and RST commands to save the DLO environment. However, it is recommended to use SAVDLO commands. When thinking of your save strategy, make sure to include DLOs into your backup procedure if you are not using the SAVE menu options.

The AS/400 system stores data in a hierarchical manner using document library objects (DLOs), which may be a document or a folder. The user's data is normally stored in a document. Documents may be collected together into a folder. Folders may contain other folders, which in turn, contain documents. Document library services maintains this hierarchical view.

In the system object view, all DLOs actually reside in libraries, and distribution (mail) documents exist in the QUSRSYS library. All other documents and folders exist in either the QDOC library or one of the QDOCnnnn libraries. The QDOC library resides in the system auxiliary storage pool (ASP). The QDOCnnnn libraries are located in the ASP corresponding to the nnnn portion of their name, ranging from 0002 to 0016. The system object name of a document or folder is a 10-character name with an internal format normally corresponding to the DLOs creation date and time.

The integrated file system provides another view of these objects using the QDLS physical file system. In this view, documents are treated as stream files. Folders are treated as directories and subdirectories.

The *Backup and Recovery*, SC41-5304, manual contains a detailed description of the commands, methods, and considerations for saving and restoring DLOs and DLO-related information. The following offers a brief overview of these elements.

15.4.1 Saving DLOs

Use the SAVDLO command to save a single document, several named DLOs, all DLOs on the system, the DLOs in a specific ASP, or only the DLOs in a specific folder. Search parameters, such as creation date or owner, may also be specified.

As an alternative method, generate a document list using the QRYDOCLIB command and the document list specified on the SAVDLO command. The advantage of this approach is that the QRYDOCLIB command provides many more search parameters than the SAVDLO command.

15.4.2 SAVDLO Enhancements

The latest enhancements to the SAVDLO command include:

- Generic name support for the FLR parameter
- OMITFLR parameter to omit some folders that match generic specifications
- Multiple, concurrent saves from a single ASP (in V4R1M0)
- Multiple, concurrent saves from different ASPs (in V3R7M0)

15.4.3 Methods of Saving Multiple Documents

The following table summarizes the SAVDLO commands to use based on the type of save that you perform.

Requirements of Save	SAVDLO Command
Single document	SAVDLO DLO(document) FLR(folder)
All DLOs	SAVDLO DLO(*ALL) FLR(*ANY)
All DLOs in a list of folders	SAVDLO DLO(*ALL) FLR(FOLDER1 FOLDER2)
Concurrent saves	SAVDLO DLO(*ALL) FLR(A* B* C*...L*) SAVDLO DLO(*ALL) FLR(M* N* O*...Z*)
All DLOs in a single ASP	SAVDLO DLO(*ALL) FLR(*ANY) ASP(n)
All DLOs in a document list	SAVDLO DLO(*DOCL) DOCL(docname)
All documents meeting certain search criteria	SAVDLO DLO(*SEARCH). See Section 15.4.4, "DLO(*SEARCH)" for a description of the possible search parameters.
All distribution objects (mail)	SAVDLO DLO(*MAIL)
All mail and all DLOs created or changed since the last complete save	SAVDLO DLO(*CHG)

15.4.4 DLO(*SEARCH)

A DLO(*SEARCH) type of search is known as a parametric search. With the parametric search, the user defines the descriptors on which to search. The user enters the information into the document details display, such as author, subject, and keywords. DLO(*SEARCH) only searches the header information of OfficeVision/400 documents. Therefore, use it as soon as OfficeVision/400 is installed and users are creating and saving documents in their folders.

Table 22 shows the options that are available if you specify DLO(*SEARCH).

Parameter	Definition
FLR	Folder
SRCHTYPE	Default is *DOC to save only documents matching search criteria. Specify SRCHTYPE(*ALL) to include folders in the search.
CHKFORMRK	Storage mark
CHKEXP	Document expiration date
CRTDATE	Creation date
DOCCLS	Document class
OWNER	Owner
REFCHGDATE	Last changed date
REFCHGTIME	Last changed time

15.4.5 Authority for SAVDLO Commands

This list summarizes the SAVDLO commands that require *ALLOBJ or *SAVSYS special authority:

- DLO(*ALL) FLR(*ANY)
- DLO(*CHG)
- DLO(*MAIL)
- DLO(*SEARCH) OWNER(*ALL)
- DLO(*SEARCH) OWNER(user-profile-name)

The user-profile name refers to a user profile different than the currently signed-on user profile.

15.4.6 Saving Office Services Information

Office services information includes database files, distribution objects, and DLOs. You can save them as demonstrated in the following figure.

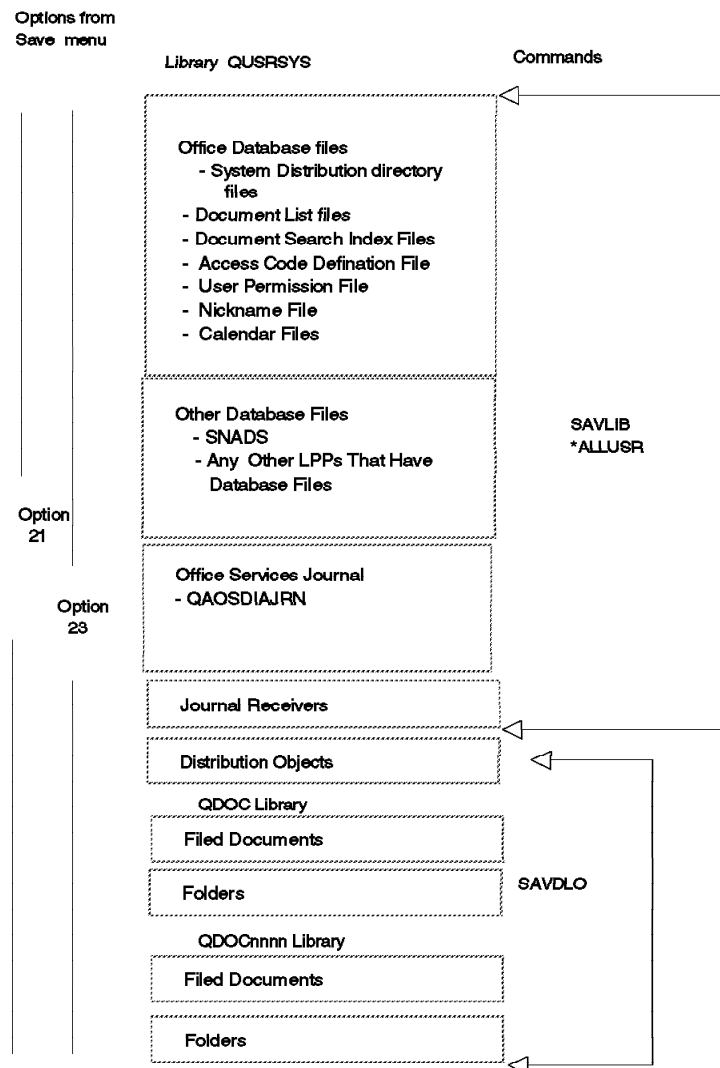


Figure 80. Saving Office Services Objects

Tips:

1. To save the complete set of Office Services information, save all documents, all mail, and the QUSRSYS library.
2. To ensure that system directory files in QUSRSYS are saved, end the QSNADS subsystem first. System directory files are not saved if QSNADS is active.

15.4.7 Saving Mail

Use the following commands to save mail:

```
SAVDLO DLO(*ALL) FLR(*ANY)
SAVDLO DLO(*CHG)
```

Use the following command to save all mail, not only changed mail:

```
SAVDLO DLO(*MAIL)
```

Remember when saving mail:

1. *ALLOBJ or *SAVSYS special authority is required to save mail.
2. Mail is dynamic in nature, so save it frequently.
3. You cannot save mail to a previous release.
4. You cannot save mail for only one user.

15.4.8 Saving Text Search Services Files

Text search services is an optional document text search tool that can be included when installing OfficeVision/400.

Text search complements the document search function available with OfficeVision/400. The document search function provided with OfficeVision/400 locates documents according to their details, such as subject, author, or date. Text search locates documents that have a specific phrase, word, word part, or combination of words or phrases in the text of the document.

The following items are saved when you save text search files:

- The administration table
- The scheduling queue
- The primary index files and additional index tables

Refer to the *Office Services Concepts and Programmers Guide*, SH21-0703, for more information on saving text search services files.

15.4.9 Restoring DLOs

Use the RSTDLO command to restore documents, folders, and mail objects. Knowing how the data was first saved determines how to restore it. This section covers restoring specific types of information using the RSTDLO command. See *Backup and Recovery*, SC41-5304, for more information.

The RSTDLO command allows you to restore:

- All saved DLOs
- Up to 300 specified DLOs
- All DLOs saved from a specified ASP
- All DLOs saved from a specified folder
- All DLOs saved at a specified date and time

If you save DLOs from more than one ASP, you must specify SAVASP(*ANY)—and the SEQNBR—which is described later in this chapter.

Note: When saving DLOs from multiple ASPs, multiple tape files are created. When you perform a SAVDLO(*ALL), a separate tape file is created for each user ASP that has DLOs.

When you specify SEQNBR(*SEARCH) on the RSTDLO command, the system restores only from the first tape file that contains DLOs. The other tape files are *not* searched unless you indicate the starting and ending sequence number on the SEQNBR parameter. Therefore, not all the DLOs are restored. If you only have DLOs in the system ASP, one tape file using the default *SEARCH value for the SEQNBR parameter works. A simple way to find the sequence numbers you need is to enter:

```
DSPTAP DEV(tape-device-name) DATA(*SAVRST) OUTPUT(*PRINT)
        ENDOPT(*REWIND)
```

This command generates a spooled file, which contains the sequence number information to use on the RSTDLO command's SEQNBR parameter.

BRMS/400 keeps track of the sequence numbers and tape files for you, so you can more easily restore DLOs from multiple ASPs.

15.4.10 Restoring New and Existing DLOs

The AS/400 system considers a restored DLO as “new” if one of these conditions is true:

- It was previously deleted.
- It is restored to a new system.
- It is restored with the NEWOBJ(*NEW) parameter. This also assigns a new 10-character, system object name.

When restoring existing DLOs, the system skips a DLO and continues if one of these conditions is true:

- The existing DLO is in use.
- The user has insufficient authority to the existing DLO.

15.4.11 RSTDLO Enhancements

The latest enhancements to the RSTDLO command include:

- Multiple, concurrent RSTDLO commands to different ASPs (in V3R7)
- Multiple, concurrent RSTDLO commands from a single ASP (in V4R2)

In V4R2, the concurrent RSTDLO process is enhanced so you can restore DLOs to the same ASP. These changes provide more flexibility for managing the DLO environment.

15.4.12 General Performance Considerations for DLOs

Restore performance is better if you have *SAVSYS special authority. Using a faster tape drive has little effect on SAVDLO performance. It takes longer to restore a DLO than to save it. It is more difficult to link the parts of a document or folder into an existing hierarchy and verify the correctness of those links. This is in contrast to moving an object already in that hierarchy to the save medium.

SAVDLO is perceived much slower than the SAVOBJ or SAVLIB command. When you save DLOs, you usually work with a large number of small objects. There is a lot of time expended on pre- and post-processing. Pre- and post-processing are performed for each and every object. When you have a lot of DLOs, you see the cumulative affect of the slower performance. All systems need sufficient memory to allow the most efficient use of time.

Concurrent saves help reduce the time to save. All systems should have sufficient memory to allow the most efficient save times.

To speed up performance, you can improve the time to save by using multiple user ASPs. For example, if it takes four hours to back up all DLOs in the system ASP, spreading the DLOs across four ASPs enables you to have a one-hour save, four times a week. The limiting factor is the amount of memory and the need to page all object headers into memory.

15.4.13 Further Considerations

When you restore DLOs, the system updates the search index database information for those DLOs. If you receive messages during the restore procedure that indicates that the information in the database does not match the DLOs, use the RCLDLO command. The RCLDLO command not only reclaims a folder, a document, or internal document library system objects, but also synchronizes the search index database with the document library.

15.4.14 Authority for RSTDLO

If you use an authorization list to secure documents and folders, specify *ALL on the Allow Object Differences (ALWOBJDIF) command, otherwise the documents are restored with *EXCLUDE authority.

<i>Table 23. Required Authorities for the RSTDLO Command</i>		
Referenced Object	Authority for Object	Authority for Library
DLO to be replaced	*ALL	*EXECUTE
Note: If you have *SAVSYS or *ALLOBJ special authority, you do not need to specify the above authority.		
Parent folder of new DLO	*CHANGE	*EXECUTE
Note: If you have *SAVSYS or *ALLOBJ special authority, you do not need to specify the above authority.		
Owning user profile of new DLO	*ADD	*EXECUTE
Note: If you have *SAVSYS or *ALLOBJ special authority, you do not need to specify the above authority.		
Tape or diskette unit	*USE	*EXECUTE
Save file	*USE	*EXECUTE
Outfile (if specified)	*USE	*EXECUTE
Note: See General Rules for Object Authorities on Commands in Appendix D-2 of the <i>Security Reference V4R1</i> , SC41-5302.		

15.4.15 Restoring Folders

When restoring folders, consider the following factors:

- You can restore a folder within a folder without restoring the one that contains the original folder.
- The fully-qualified path name specified on the RSTFLR parameter must exist unless you restore a first-level folder. You can create folders along this qualified path if necessary.
- If you restore a folder over an existing folder that is in use, the folder and all DLOs in it are bypassed and not restored.
- If you restore objects into an existing, damaged folder and you cannot reclaim it, the folder and all DLOs in it are bypassed and not restored.

15.4.16 Restoring Mail and Distribution Objects

OfficeVision/400 manages AS/400 mail through internal objects stored within the QUSRSYS library. The SAVDLO and RSTDLO commands save and restore these internal objects.

Consider the following guidelines to restore mail and distribution objects:

- Use RSTDLO DLO(*MAIL) to restore only:
 - Filed distribution objects with a mail log-reference at the time they are saved
 - All other distribution objects
 - All distribution documents
- Use RSTDLO DLO(*ALL) SAVFLR(*ANY) to restore:
 - All objects restored by RSTDLO DLO(*MAIL)
 - All documents
 - All folders
- If you used SAVDLO DLO(*MAIL) to save the mail, specify RSTDLO DLO(*ALL) SAVFLR(*ANY) to restore mail.
- You cannot restore distribution documents and objects individually.
- Mail-log references are updated for local recipients of restored documents.
- Mail-log references on remote systems for remote recipients are not restored.
- Mail-log references are restored for a local sender of a document if an entry was in the sender's mail log at the time the distribution was saved.
- Entries in the mail logs of remote senders are not saved or restored.

15.4.17 Authority and Ownership Issues a During a Restore of DLOs

Here is a summary of authority and ownership rules when restoring DLOs.

- Rule 1** If the owning user profile of the saved object exists on the system, that user profile owns the restored object.
- Rule 2** If the owning user profile of the saved object does not exist on the system, QDFTOWN (system default owner) owns the restored object.
- Rule 3** If the object exists on the system and is owned by a different user profile than the owner of the saved object, the object is not restored unless ALWOBJDIF(*ALL) is specified. In that case, the object is

owned by the user profile of the existing object, not the owner of the saved object.

Rule 4 If the owning user profile of the saved object is not enrolled in the system distribution directory ownership, QDFTOWN owns the restored object.

Rule 5 When restoring a new DLO to the system, all access codes or private user authorities are removed. The private authorities to DLOs are restored upon completing the RSTAUT command. However, the access codes are not restored.

15.4.18 Recovery of Text Index Files for Text Search Services

The text index database files are saved when library QUSRSYS is saved. If you restore the text index files, restore all files together from the same backup. Otherwise, their association with each other is lost, which causes unpredictable results. If you do not have saved copies of the files, restore them from the original distribution tapes.

If you recover all text index search files and documents from the same set of save tapes, you should not have problems. If you recover the system in pieces, consult Appendix F, Procedures for Recovering the Text Index, in the *Backup and Recovery* manual, SC41-5304.

15.5 Domino for AS/400

Domino for AS/400 runs on OS/400 similar to other AS/400 applications. Because of this relationship, significant improvements are realized in scalability, integration, and management capabilities.

Running Domino natively maximizes the integrated access to the AS/400 system and its data by taking advantage of the AS/400 integrated file system. This implementation also offers major benefits of performance and scalability. The reliability of the AS/400 system offers a significant advantage over other Domino options. Domino for AS/400 is managed like all other AS/400 applications using the same resources and skills. It fully exploits the 64-bit server technology of the AS/400 system for added performance.

For availability, Domino for AS/400 has an automatic restart capability. If the AS/400 system detects that the Domino server has unexpectedly stopped, it automatically brings the server back online.

Note: Domino for AS/400 is available only on RISC-based processors (running OS/400 V4R2 or later), while both RISC and CISC systems can use the integrated PC Server implementation.

15.5.1 Why You Should Back Up a Domino for AS/400 Server

A typical Domino server is a source of business information. This information might not exist anywhere else other than on your servers. Typical scenarios include:

- Users relying on e-mail for communication that is not duplicated anywhere else
- An online customer service application that contains records that do not exist on hardcopy

In developing a backup strategy, consider two different object categories:

- Objects that change infrequently, such as programs for the Domino product. See Section 15.5.3, “Backing Up the Domino for AS/400 Product,” for more information.
- Objects that change frequently or regularly, such as Domino databases on the server. See Section 15.5.4, “Backing Up the Domino for AS/400 Server” on page 250, for more information.

To understand how to back up the various components, you need to understand what the libraries and directories are for this product. The AS/400 system has the concept of a single-level storage. In single-level storage architecture, the system manages the allocation of disk space. When you back up information, you back up logically (by library or directory), not physically (by disk units).

15.5.2 Libraries and Directories for the Domino for AS/400 Product

The following two tables list the libraries and directories containing Domino related information.

Description	Library	Integrated File System Paths for Library
Domino for AS/400 Product	QNOTES	/QSYS.LIB/QNOTES.LIB
Directory synchronization	QNOTESINT	/QSYS.LIB/QNOTESINT.LIB
C APIs	QNOTESAPI	/QSYS.LIB/QNOTESAPI.LIB
C++ APIs	QNOTESCPP	/QSYS.LIB/QNOTESCPP.LIB
HiTest APIs	QNOTESHTST	/QSYS.LIB/QNOTESHTST.LIB
LotusScript Extensions	QNOTESLSKT	/QSYS.LIB/QNOTESLSKT.LIB
Customization information (SBSDs and JOBDs)	QUSRNOTES	/QSYS.LIB/QUSRNOTES.LIB

Description	Integrated File System Path for Directories
Product information	/QIBM/ProdData/LOTUS/NOTES
Customization files	/QIBM/UserData/LOTUS/NOTES
Directory for databases on the server	Specified when you configure the server

15.5.3 Backing Up the Domino for AS/400 Product

We recommend that you back up Domino for AS/400 after you install the product. Also, save the product after applying fixes to it. Otherwise, the information in these products is considered relatively static.

To save the static information, use these options:

1. Use option 21 from the SAVE menu to save the entire system.
2. Use option 22 from the SAVE menu to save system data only. Option 22 saves the product libraries and directories, including the QNOTESxx libraries and the /QIBM/ProdData/LOTUS/NOTES directory.

15.5.4 Backing Up the Domino for AS/400 Server

When saving a Domino server, you typically save all the dynamic information associated with the server.

You must save:

- All the databases
- Users' mail databases
- The names and address book for the server

When you configure a Domino server, specify the directory for that server such as /NOTES/DATA. By default, all the databases for the server are in that path. Typically, end users cannot create Domino databases in any location except the default path for the server. If you are responsible for backing up a Domino for AS/400 server, develop a backup strategy that matches your policy of where you keep information. The two main approaches for this are:

- Limit the location of Domino databases.
- Save everything.

15.5.4.1 Limiting the Location of Domino Databases

By limiting the location of Domino databases, you need to back up only the server path and the path that contains IBM-supplied dynamic objects (/QIBM/UserData/Lotus/Notes). Also, use a combination of policies and security to keep all Domino databases within the default directory (path) for the Domino server.

The following example outlines the steps to back up the directory for your Domino for AS/400 server and the IBM-supplied directory:

1. Sign on to the AS/400 system as a user with *JOBCTL and *SAVSYS authority.
2. End the Domino server before you proceed. This ensures that the save is complete. The command to end the server is:

```
ENDDOMSVR SERVER(server-name)
```

3. To save your Domino directory and the system-supplied directory, enter the command:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD')  
OBJ('/Notes/data/*')  
  ('/QIBM/UserData/Lotus/Notes/*')
```

Substitute the /NOTES/DATA with your directory name.

15.5.4.2 Save Everything

By saving everything, you assume that Domino databases exist anywhere in the integrated file system (in either the root directory or the QOpenSys directory). Therefore, back up the entire root directory and the /QOpenSys directory. Use any of these methods:

1. Option 21 from the SAVE menu to save the entire system
2. Option 23 from the SAVE menu to save all user data
3. SAV command to save everything except QSYS.LIB, the /ProdData directory, and the QDLS file system:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD) OBJ('//*') ('QSYS.LIB' *OMIT)
('/QIBM/ProdData' *OMIT)
('/QDLS' *OMIT) (QOpenSys/QIBM/ProdData' *OMIT))
UPDHST(*YES)
```

Note: When using options 21 or 23 from the SAVE menu, the system automatically ends all other activities on the system and enters the restricted state. However, if you use commands to save the server information, stop the server before the saves. To stop the server, issue the End Domain Server (ENDDOMSRV) command:

```
ENDDOMSRV SERVER(server-name)
```

15.5.5 Backing Up Specific Dynamic Objects From Your Domino Server

You can save objects from the Domino server by the type of object. The following section describes how to back up specific dynamic objects from the Domino server.

15.5.5.1 Backing Up Mail from Your Domino Server

If you already have a strategy for saving all the user information from your Domino server, you probably do not need a separate procedure for saving only mail. If your backup interval for your entire server is infrequent, consider saving mail because mail objects are volatile.

Your Domino server stores mail in multiple databases:

- The MAIL.BOX database on each server contains mail for the server to route to individual user mailboxes or to another server.
- Each Lotus Notes user has an individual mail database, which is typically the user's ID with the NSF extension. All individual mail databases are in a dedicated subdirectory such as /NOTES/DATA/MAIL.

To save mail, perform the following steps:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. To save a database, the AS/400 system must lock the database so that no changes occur during the SAV operation. To successfully save NOTES mail, stop the Domino for AS/400 server that contains the mail databases.

To stop the Domino server, enter:

```
ENDDOMSRV SERVER(server-name)
```

3. Mount the media into the tape drive.
4. Enter the command:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD')
OBJ('/NOTES/DATA/MAIL/*.NSF')
```

15.5.5.2 Examples of Backing Up Mail from Your Domino Server

The following table shows some examples for backing up mail objects.

<i>Table 26 (Page 1 of 2). Examples of Backing Up Mail From Your Domino Server</i>	
Commands	Description
OBJ('/NOTES/DATA/MAIL.BOX')	Saving a specific database such as MAIL.BOX
OBJ('/NOTES/DATA/MAIL/*.NSF')	Saving files of a specific type in the MAIL subdirectory

<i>Table 26 (Page 2 of 2). Examples of Backing Up Mail From Your Domino Server</i>	
Commands	Description
OBJ('/NOTES/DATA/MAIL/DMT.NSF')	Saving a specific user's mail database such as DMT mail database
Note: We assume that the directory for your Domino server is /NOTES/DATA.	

15.5.5.3 Backing Up a Specific Database

To create a backup copy prior to programming changes or to test a new agent, you may want to back up a specific database. You may also want to back up the database to create an archived copy at the end of an accounting period.

For example, to back up a specific database called CUSTINF.NSF:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. Ensure that no one is using the database.
3. Enter in the SAV command:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD')
OBJ('/NOTES/DATA/CUSTINF.NSF')
```

The following table shows examples of the SAV commands that you can use to back up specific Domino databases.

<i>Table 27. Examples of Backing Up Specific Domino Databases</i>	
Commands	Description
OBJ('/NOTES/DATA/DEPT57/*.NSF')	Saving all the Domino databases in the DEPT57 subdirectory
OBJ('/NOTES/DATA/HRDPT/HRINFO.NSF')	Saving HRINFO database from the HRDPT directory
OBJ('/NOTES/DATA/HR*.NSF') Note: We assume that the directory for your Domino server is /NOTES/DATA.	Saving all the HR files that begin with "HR"

Note: Before the save operation, verify that no one is using the database, or simply end the server.

15.5.5.4 Backing Up Changed Objects from Your Domino for AS/400 Server

Consider the example of an operation to save all objects in the HRINFO directory that changed since 8:00 A.M. on November 8, 1997. Enter the SAV command:

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD')
OBJ('/NOTES/DATA/HRINFO/*.*)
CHGPERIOD('11/08/97' 080000)
UPDHST(*YES)
```

Here are two scenarios for saving changed objects from your Domino for AS/400 server:

1. Back up all changes since the last full system backup, as shown in the following figure.

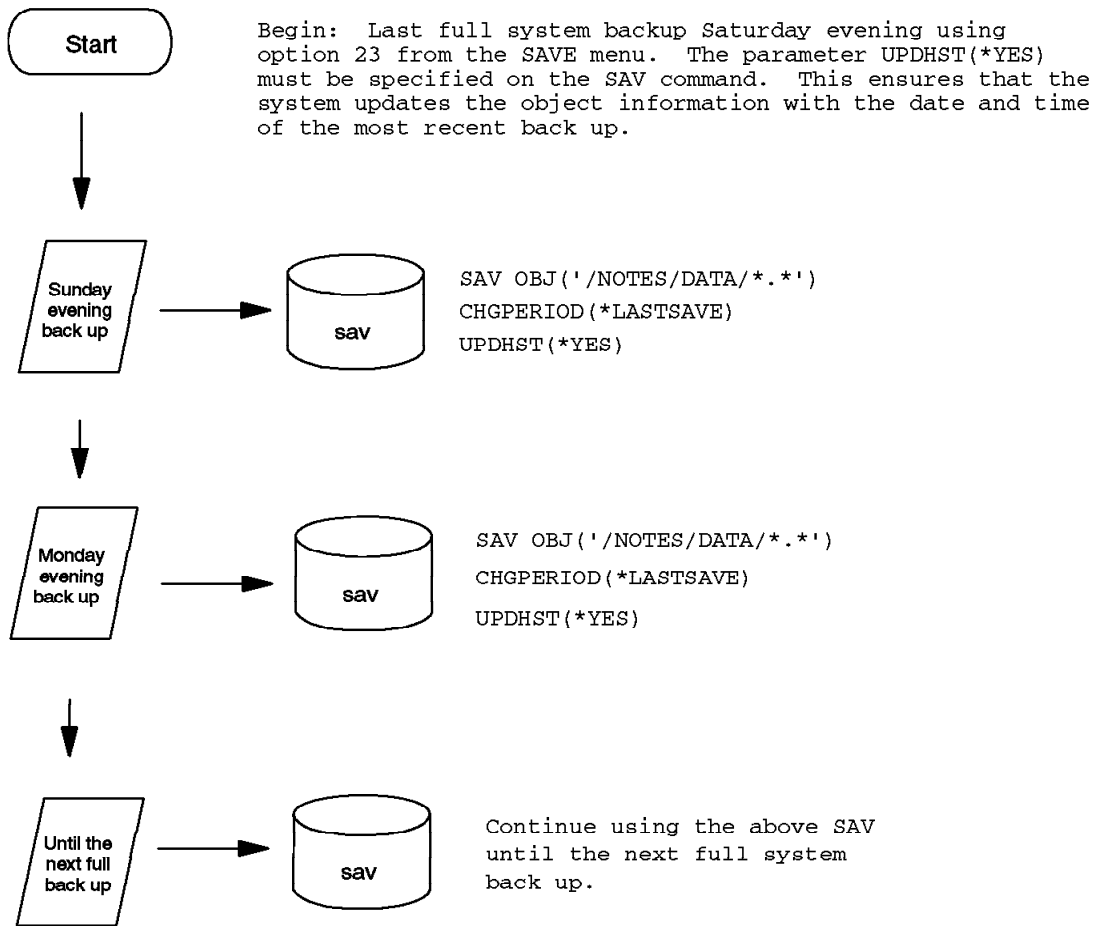


Figure 81. Scenario 1 — Saving Objects Changed Since Last Full BackUp

2. Back up daily changes, as demonstrated in the following figure:

Note: This example assumes that your Domino server directory is in /NOTES/DATA. We also left out the DEV parameter. This is where you need to specify the tape device to use (for example, DEV('/ QSYS.LIB/tape-device-name.DEVD')).

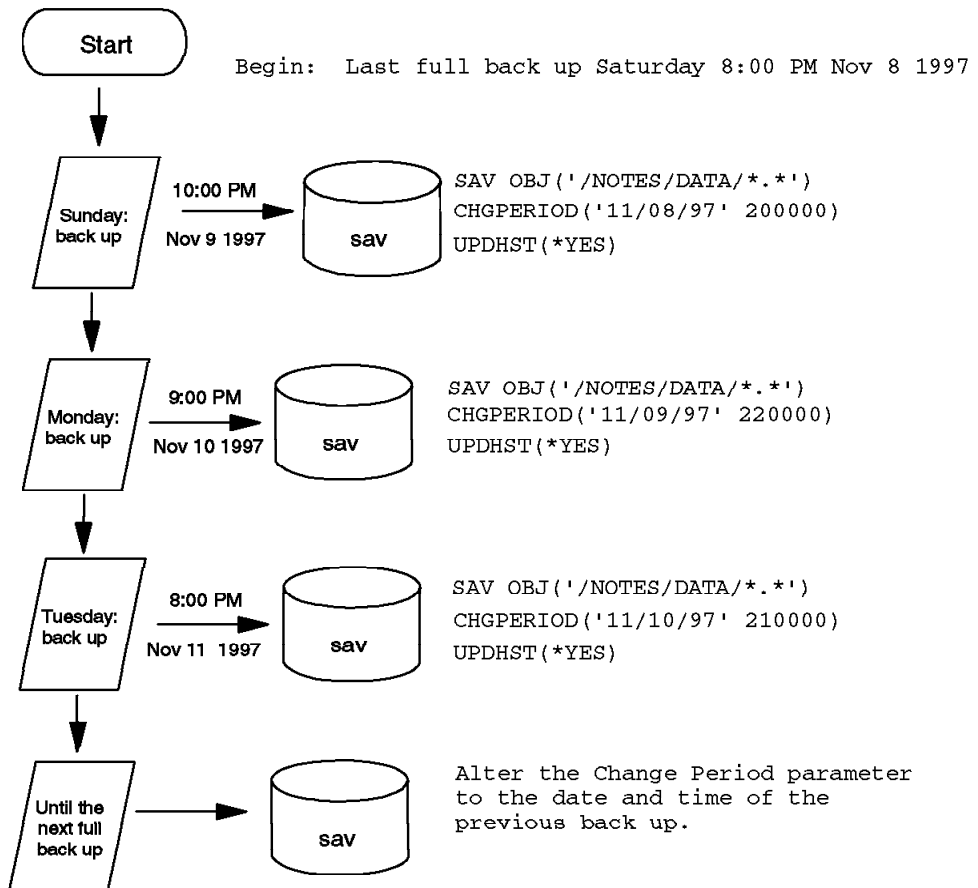


Figure 82. Scenario 2 — Saving Changed Objects Daily

15.5.6 Recovery of Domino for AS/400

The Domino for AS/400 product resides in libraries in the QSYS.LIB file system on your AS/400 system. Table 24 on page 249, in the previous section on saving Domino for AS/400, describes the Domino for AS/400 libraries in the QSYS.LIB file system.

The next few sections describe recovery steps for the Domino for AS/400. Consult the *Backup and Recovery* manual, SC41-5304, and the Domino documentation to better understand Domino recovery. You must restore objects in the correct sequence to rebuild the proper links between objects.

15.5.6.1 Recovering an Entire Domino for AS/400 Server

Follow these steps as an example of how to recover the entire Domino for AS/400 server:

1. Sign on to the AS/400 system with a user profile that has *JOBCTL and *SAVSYS authority.
2. End the Domino server. This ensures that no one is using the server that you plan to restore. Enter:
ENDDOMSRV SERVER(server-name)

3. Mount the tape containing the most recent backup copy of the directories for the server.
4. Restore your Domino directory using the command (assuming your Domino directory is /NOTES/DATA):

```
RST DEV('/ QSYS.LIB/tape-device-name.DEVD')
OBJ('/ NOTES/DATA/*')
```

This process restores the physical contents of your server to your AS/400 system. Consult the Domino documentation for any special recovery activities that you need to perform after restoring the directories.

15.5.7 Recovering Domino Mail

If you need to recover one or more mail databases from backup media, use the Restore (RST) command. Follow these steps:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. Stop the server containing the mail databases you want to restore. Issue the command:

```
ENDDOMSRV SERVER(server-name)
```

3. Mount the tape containing the most recent backup copy of database.

To restore all the databases to the MAIL subdirectory, for example, issue the command:

```
RST DEV('/ QSYS.LIB/tape-device-name.DEVD')
OBJ('/ NOTES/DATA/MAIL/*')
```

15.5.7.1 Examples of Recovering Mail from Your Domino Server

The following table offers examples for recovering mail.

Commands	Description
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/NOTES/DATA/MAIL/DMT.NSF')	RST DEV('/QSYS.LIB/tape-device-name.DEVD') Restoring a specific user's mail database, such as DMT's mail
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ((' /NOTES/DATA/MAIL/ASHUSR.NSF') ('/NOTES/DATA/MAIL/ALIUSR.NSF') ('/NOTES/DATA/ADRIUSR.NSF'))	Restoring mail databases for ASHUSR, ALIUSR, and ADRIUSR

Notes:

1. You cannot restore over a database that is in use. All users must close the database before you can restore a backup copy. To determine if a database object is in use, use the *WRKOBJLCK* command.
2. All examples above assume that the directory for your Domino server is /NOTES/DATA.

15.5.8 Recovering a Specific Database

In some special cases, you may want to restore a specific database or a group of databases. Use the restore (RST) command to restore a specific database. To recover a specific database, follow the steps of this example:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. End the server containing the database you want to restore.
ENDOMSRV SERVER(server-name)
3. Mount the tape containing the most recent copy of the database.

To restore all the files to the MAIL subdirectory, called HRDPT for example, issue the command:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/NOTES/DATA/HRDPT/*.NSF')
```

15.5.8.1 Examples of Recovering the Domino Databases

The following table contains some examples of recovering Domino databases.

Commands	Description
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/NOTES/DATA/HRDPT/HRINFO.NSF')	Restoring a specific database HRINFO to the HRDPT subdirectory
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/NOTES/DATA/CUSTSVC/*.NSF')	Restoring all the Domino databases to the CUSTSVC subdirectory
RST DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/NOTES/DATA/DMT*.NSF')	Restoring all the Domino databases beginning with DMT to the main subdirectory of your server

Notes:

1. You cannot restore over a database that is in use. All users must close the database before you can restore a backup copy.
2. All examples that are shown assume that the directory for your Domino server is /NOTES/DATA.
3. In the DEV parameter, enter the media device from which you want to restore, such as:

```
DEV('/QSYS.LIB/tape-device-name.DEVD')
```

15.5.9 Restoring Changed Objects to the Domino for AS/400 Server

If your save strategy includes saving changed objects, plan your recovery sequence carefully. We present four scenarios:

1. Restoring changed objects from a cumulative backup
2. Restoring changed objects from nightly backup
3. Restoring a Domino database from incremental backup
4. Restoring changed objects to a specific Domino subdirectory

15.5.9.1 Restoring Changed Objects from a Cumulative Backup

If your strategy for saving changed objects is cumulative, that is you save everything that was changed each night, follow these recovery steps:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. End the Domino server to ensure that no one is using the databases. Enter:

```
ENDDOMSRV SERVER(server-name)
```

3. Mount the tape containing the most recent *complete* backup copy of the databases.

4. Use the RST command to restore the entire Domino database directory:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')  
OBJ('/NOTES/DATA/*')
```

5. Mount the most recent save tapes containing your changed objects.
6. Restore all the objects changed since the last full back up, by issuing the command:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')  
OBJ('/NOTES/DATA/*')
```

15.5.9.2 Restoring Changed Objects from Nightly Backup

If you save changed objects every night or every day, you only save objects that change since the previous night or day. Follow these recovery steps:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. Ensure that no one is using the databases. Enter:

```
ENDDOMSRV SERVER(server-name)
```

3. Mount the tapes from the most recent *complete* back up.

4. Use the RST command to restore the entire Domino database directory:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')  
OBJ('/NOTES/DATA/*')
```

5. Mount the *first* save tapes containing the changed objects. For example, if your last save of *everything* was on Saturday night, locate your save tapes from Sunday night.

6. To restore the objects on the tape (that is everything that changed since the previous night), issue the command:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')  
OBJ('/NOTES/DATA/*')
```

7. Repeat steps 5 and 6 for each save tape until you attain currency in your directory.

15.5.9.3 Restoring Domino Database from Incremental Backup

The steps to restore a specific database named HRINFO to the HRDPT subdirectory are:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.

2. Ensure that no one is using the databases. Enter:

```
ENDDOMSRV SERVER(server-name)
```

3. To locate the most recent tape with the database on it, complete these steps:

- Scrutinize the job log of the save job.

- Use the DSPTAP command to verify the contents of the tape.
4. Mount the tape with the incremental backup on the tape drive.
 5. To restore the database, enter the command:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
OBJ('/NOTES/DATA/HRDPT/HRINFO.NSF')
```

15.5.9.4 Restoring Changed Objects to a Specific Domino Subdirectory

To restore all the Domino databases to the CUSTSVC subdirectory, complete these steps:

1. Sign on with a user profile that has *JOBCTL and *SAVSYS authority.
2. Ensure that no one is using the databases. Enter:

```
ENDDOMSRV SERVER(server-name)
```

3. Mount the tape containing the most recent *complete* backup.
4. Use the RST command to restore the entire directory from the tapes:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
OBJ('/NOTES/DATA/CUSTSVC/*')
```

5. If your incremental backups are cumulative, mount the most recent incremental backup tape. Use the same restore command (see step 4) to restore the changes.
6. Repeat steps 3 and 4 for each incremental backup tape. Start with the oldest tape and work forward using the RST command in step 4.

15.6 Windows NT

One of the advantages of the Integrated PC Server and Windows NT is the incorporation of the Windows NT backup procedures into the AS/400 system backup. A big advantage is the ability to use AS/400 tape media as a backup device.

The environment is a combination of both Windows NT and OS/400 for which there are several backup options. A complete storage space backup is accomplished with an AS/400 system save. This is the fastest method to back up the Windows NT environment, however, you cannot restore individual files using this strategy.

Note: The network server description (NWSD) of the Windows NT server must be varied off before performing OS/400 backup and recovery operations on its storage spaces.

This chapter describes information that you can use for designing a backup plan for Windows NT within the integrated file system.

15.6.1 Directories and Objects for Windows NT

Windows NT uses the directories and objects listed in the following table. The object name is derived from the network server storage description. For our example in Table 30 on page 259 the NWSD is named NTTEST.

Object Location	Object Name	Object Type	Object Content	Save Command
QUSRSYS	NTTEST1	Server Storage Space	DOS boot drive (C: drive)	SAVOBJ OBJ(NTTEST1) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG)
QUSRSYS	NTTEST2	Server Storage Space	Code copied from Windows NT CD (D: drive)	SAVOBJ OBJ(NTTEST2) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG)
QUSRSYS	NTTEST3	Server Storage Space	Windows NT install drive (E: drive)	SAVOBJ OBJ(NTTEST3) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG)
QNTAP	QNTAP	Library	AS/400-based Windows NT code	SAVLIB LIB(QNTAP) DEV(tape-device-name)
QGPL	NTTESTQ	Server Message Queue	Messages from the Windows NT server	SAVOBJ OBJ(NTTESTQ) LIB(QGPL) DEV(tape-device-name) OBJTYPE(*MSGQ)
/QIBM/ProdData	NTAP (and sub-directories)	Directory	PC-based Windows NT code	SAV DEV('QSYS.LIB/tape-device-name.DEVD') OBJ('/QIBM/ProdData/NTAP')
QSYS	NTTES* (various)	Device Configuration Objects	AS/400 devices for Windows NT	SAVCFG DEV(tape-device-name)
/QFPNWSSTG	various	User Storage Spaces	User data and applications	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QFPNWSSTG/server-storage-space-name')

Note: The server message queue is typically not saved, because messages are volatile and do not need to be restored to regain production.

Generally, we divide Windows NT product information as:

- System objects that are relatively static
- User objects that are dynamically changed

15.6.2 Backing Up System Objects

We recommend that you back up Windows NT after you install the product. Also, back it up after installing any fixes to it.

System objects are created during the installation of the AS/400 integration feature. They are considered part of the AS/400 system and are saved during a full system backup.

To save static system objects or static information, use the following options:

- Use option 21 to access the SAVE Menu, and select Save the Entire System.
- Use option 22 System Data Only from the SAVE menu, if you complete option 23 on a regular basis. Use option 22 after an installation or upgrade to the product, or after PTFs are applied.

15.6.3 Backing Up User Objects

User objects on the AS/400 system are not part of the operating system and are not required to operate the NT server. User storage spaces are considered user objects. User storage spaces exist in the integrated file system directory /QFPNWSSTG.

To save the user storage space named "DISK1" in the /QFPNWSSTG directory to the tape media device, enter the command:

```
SAV DEV('QSYS.LIB/tape-device-name.DEVD')
OBJ('/QFPNWSSTG/DISK1')
```

The RST command is used to restore.

15.6.4 Considerations for Back Up when Creating User Spaces

Plan ahead when creating user spaces to ensure that the backup is performed easily and efficiently. Follow these recommendations:

- Do not store any user data on the E: drive.
- Keep static data (such as applications and software) and frequently modified data (such as user data) on separate storage spaces.
- Keep all data belonging to one group or application on one storage space because it is easier to manage.

Applications typically do not change often so they can be omitted from a daily back up routine. Frequently modified data should be backed up frequently.

15.6.5 Backing Up Specific Objects From Windows NT

When you perform a full AS/400 system backup, any Integrated PC Servers on the AS/400 system are automatically saved during the process, provided that you use the SAVE menu options to perform the backup. If you use CL programs to perform the saves, vary off the IPCS and use the SAV OBJ('/*) command that Option 21 from the SAVE menu calls out, as:

```
SAV OBJ('/*) ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)) UPDHST(*YES)
```

The following are options and considerations for specific objects from the Windows NT system.

15.6.5.1 Windows NT Operating System and Registry

The Windows NT system and the files required to boot the Windows NT server are located on the C: , D: , and E: drives of the server. The server storage spaces that contain these drives are in the QUSRSYS library. See Table 30 on page 259, which illustrates where the drives are located.

Because of the structure of the Integrated PC Server, save these drives on the AS/400 system and restore them if files on these drives are damaged or deleted. For instance, perform the following actions:

- If your BOOT.INI is deleted, restore the C: drive.
- If the /386 on the D: drive is deleted, replace it by restoring the D: drive.
- If the Windows NT registry on the E: drive is corrupted, replace it by restoring the E: drive.

Save these objects frequently. Save the E: drive daily if possible, since it contains the registry. Then, save it to a save file.

Use the SAVOBJ command to save the Windows NT operating system and registry:

```
SAVOBJ OBJ(NTTEST3)  
LIB(QUSRSYS) DEV(*SAVF)  
OBJTYPE(*SVRSTG) SAVF(NTBACKUP/SVRSTG3)
```

15.6.5.2 User Data

The only AS/400 save command that saves Windows NT user data is the SAV command. You cannot restore individual files and directories using SAV and RST. You can only save the user storage space as a whole entity and restore the storage space as a whole entity.

To save a user space named DISK1, enter the SAV command:

```
SAV DEV(' QSYS.LIB/tape-device-name.DEVD')
    OBJ('/ QFPNWSSTG/DISK1')
```

Typically, users employ PC backup utilities to save user data.

15.6.5.3 Backup to AS/400 Tape

The AS/400 tape device can be logically detached from the AS/400 system and reassigned to the Windows NT system. Windows NT sees the tape device as an attached, physical tape device. This enables Windows NT to use AS/400 tape devices.

AS/400 tape devices supported from the Windows NT IPCS as native NT devices include:

- 6385 13G 1/4" Cartridge Tape Unit
- 3590 1/2" Cartridge High Performance Tape Subsystem
- 3570 8mm Tape Cassette Subsystem
- 1/4" cartridge tape devices with device type of 63A0 and 6385
- 3494 models L1 and D12
- 3570 models B00, B01, B11 and B1A
- 3570e models C00, C01, C11 and C1A
- 3590 models B11 and B1A
- 3590e models B21 and B2A
- 6381
- 6382
- 6385
- 6386
- 6390
- 7208 (all models)
- 9427 models 210, 211, 310 and 311

Refer to APAR II11119 for additional information on supported tape drives for the Windows NT environment. See *Basic System Operation, Administration, and Problem Handling*, SC41-5206-01, for information on accessing APARs.

Once the tape is attached to the Windows NT system, use it as you would use a PC-based tape device. Use one of the supported backup utilities, such as the Windows NT backup utility or the Seagate "Backup Exec" to direct your backups to the AS/400 tape device. When using the Windows NT tape support, the AS/400 tape drive must be varied off, and therefore, is not available for other tape operations from AS/400 commands. The tape drive can be used by the Windows NT server, and before using a Windows NT backup product once the tape is varied off and is locked for use. Backups from AS/400 save operations cannot be mixed on tapes containing backups from PC utilities.

IBM's ADSTAR Distributed Storage Manager (ADSM) can also be used for an unattended network backup and archive service facility at a file level.

15.6.5.4 Copying the Files to a Directory on the AS/400 System

Another option for saving individual files and directories is to save to the AS/400 integrated file system.

There are several ways to move the files from the Windows NT server into the integrated file system.

- OS/400 Support for Windows Network Neighborhood (AS/400 NetServer)

For V4R2 systems there is an enhancement called NetServer that uses server message block (SMB) to allow the integrated file system of the AS/400 system to be visible to the clients on the network without additional software on the client. NetServer is can be used to copy files from Windows NT or Windows 95 to the integrated file system.

- AS/400 Client Access

You can install AS/400 Client Access on the Windows NT and Windows 95 server running with the Integrated PC Server or AS/400 system. Once installed, use its access to the AS/400 integrated file system to copy files and directories there.

To copy files from Windows NT to the integrated file system, open the network neighborhood on the Windows NT server, and select the AS/400 option and the location within the integrated file system to where you want to copy the files. Drag and drop the files as you usually do. Reverse the process to restore the files.

- Other backup utilities

There are many third-party utilities available that enable you to back up your files on the network, such as WinTAR-Remote.

15.6.6 Restoring the Windows NT Product

Basically, use the same parameters on the restore command that you use on the save command. You must follow the restore sequence.

Before restoring components (using Windows NT utilities), you must have an operational Windows NT server. If you lose the server, restore the server storage spaces and start the server up. Note these important points:

- Windows NT server system files

If you accidentally destroy some of the system files, you can recreate them from the D: drive. The Windows NT structure uses a DOS shell-based expand program that allows users to restore files that are kept as images on the D: drive. Use the EXPAND command to expand the needed files from the D:\386 directory back to the related directory on the E: drive. For example, to expand MFC40U.DLL to the appropriate directory on the E: drive, enter:

```
D:\386\EXPAND D:\386\MFC40U.DL_ E:\386\SYSTEM32\MFC40U.DLL
```

Restore critical Windows NT server system files (such as NTOSKRNL.EXE and WIN32K.SYS). These are files that Windows NT needs to function properly. It is possible to restore these and other files from an image stored on the D: drive to the E: drive. You can only use files that are already expanded because of the lack of a true DOS-based EXPAND program. If the file exists on the D: drive in expanded form, use the COPY command to replicate the file into the appropriate directory on the E: drive.

Note: To find out whether a file exists in expanded form, check the three-character extension on the desired name in the D: drive. If the extension contains an underscore character (_), the file is compressed on the D: drive. If the file exists on the D: drive in an expanded form, the proper three character extension appears in the filename. For example, for WIN32K.SYS, the compressed form is WIN32K.SY_.

You can restore Windows NT server AS/400 device drivers that are accidentally destroyed. An expanded copy of all these device drivers is located in the D: drive in the directory: D:\386\%oem%. Copy the appropriate file from the D: drive back to the E: drive.

- Windows NT operating system and registry

To restore the E: drive, vary off the Integrated PC Server, restore storage space three (example NTTEST3) to the QUSRSYS library, and vary on the Integrated PC Server again. To complete the restore operation, enter the command:

```
RSTOBJ OBJ(NTTEST3)
        SAVLIB(QUSRSYS) DEV(*SAVF)
        OBJTYPE(*SVRSTG) SAVF(SVRSTG3/NTBACKUP)
```

Restore the C: and D: drives (storage spaces one and two respectively) in a similar way.

Note: The registry information is stored on the E: drive. It contains references to the applications installed on the server, and on the domain control database. Be careful because you can restore a back-level version of a backup over the current one.

- User storage spaces

If you save a storage space from /QFPNWSSTG directory, only restore the whole storage space. You cannot restore individual files from this backup.

- File level restore

After restoring the user storage spaces, check when they were saved. Restore any incremental backups from this point.

Note: If you restore the configuration objects using the RSTCFG command, you *must* run the RSTCFG command before restoring any of the server storage spaces. When RSTCFG restores the configuration objects for the server storage spaces, it runs a similar process to the INSWNTSVR command that initializes the storage spaces on the AS/400 system. If existing server storage spaces are found, the restore of the object and the restore of the network server description both fail.

- Restoring a domain controller

Be careful when you restore either a Primary Domain Controller (PDC) or a Backup Domain Controller (BDC). Ensure that the domain database held on the server is synchronized with the other domain controllers. This applies particularly if the server you restore was previously the PDC that was replaced in that role by one of the BDCs in the domain.

In the Integrated PC Server, normal Windows NT procedures apply. Be aware of any changes in your current setup of PDC or BDCs since the last backup.

- Re-install the server

Re-installing the server is done as a last resort only.

For a complete description of backing up and restoring Windows NT, refer to *AS/400--Implementing Windows NT on the Integrated PC Server*, SG24-2164.

15.7 Lotus Notes on the Integrated PC Server

Lotus Notes is the leading groupware platform in the industry today. Groupware, as exemplified by Lotus Notes, enables and encourages people to work in teams. It allows them to share, process, and view information in a way that is useful to them and their business. Key features of Lotus Notes include:

- Tools and interfaces so users can easily create and share information
- Support for sending, receiving, and managing electronic mail
- An environment for creating and running Notes applications customized to suit business needs

Notes is designed for the client/server environment. A typical Notes environment consists of client (user) workstations that communicate with one or more Notes servers through a network.

The server contains databases of information that workstation users can share and jointly update. One special database is a mailbox for electronic mail. Another special database is an address book that provides directory services.

15.7.1 Types of Storage Spaces

An Integrated PC Server containing a Notes server has at least six storage spaces. Each of these storage spaces is represented by a drive letter— C:, D:, E:, F:, G:, and K:.

Note: These storage spaces are backed up when you enter the SAVLIB LIB(QUSRSYS) command, or you enter SAVLIB LIB(*NONSYS or *ALLUSR).

- C: drive

When you create a network server description, you automatically create the C: drive as one of two server storage spaces. These drives have files, such as CONFIG.SYS and STARTUP.CMD, that are needed to boot the OS/2 Warp operating system for the network server. The storage space is named xxxx1, where xxxx is the name of the network server. For example, if the network server is named SERVER1, the storage space is named SERVER11. This is in library QUSRSYS. You cannot change the files in this storage area.

- D: drive

The D: drive contains the OS/2 Warp operating system that runs the Integrated PC Server. The AS/400 system creates this storage space when you install Integration Services for FSIOP. This storage space is named QFPBSYS2. This is in library QFPINT. You cannot change the files in this storage space.

- E: drive

When you create a network server description for the Integrated PC Server, you create the E: drive, which is the second of the two server storage spaces automatically created. Each network server has its own E: drive, which contains start-up and log files for the network server.

The E: drive storage space is named xxxx3, where xxxx is the name of the network server for which this storage space was created. For example, if the

network server is named SERVER1, the storage space is named SERVER13. The object resides in library QUSRSYS.

If you change a file in the storage space, save the storage space.

- F: drive

The F: drive contains programs and data that are used by OS/400 Integration of Lotus Notes. These programs and data are shared by all network servers using the Notes integration functions. This storage space is in the QUSRSYS library, and is named QYNASYS1.

You cannot change the files in this storage area.

- G: drive

The G: drive stores Notes source images (compressed files) that you use later to set up either a Notes server or related Notes functions. This drive is divided into two segments that are identified as location 1 and location 2.

This storage space is in the QUSRSYS library. The storage space is named QYNALNCD.

- K: drive

The K: drive contains Notes server files and Notes databases supported by the Notes server on the Integrated PC Server. The K: drive is the network server storage space created when you set up the Integrated PC Server to receive the Notes server. When creating this storage space, you have to specify the size and name.

To add files to or remove files from the Notes server, you must access this storage space.

Note: These storage spaces are backed up when you save the QUSRSYS library using one of the following commands:

```
SAVLIB LIB(QUSRSYS)
SAVLIB LIB(*NONSYS)
SAVLIB LIB(*ALLUSR)
```

The following figure shows the types of storage spaces on the Integrated PC Server.




Figure 83. Types of Storage Spaces

This information is also shown in following table.

<i>Table 31. Table Showing Libraries and the Objects for Notes on IPCS</i>				
Object Location	Object Name	Object Type	Object Content	Save Command
QUSRSYS	SERVER11	Server Storage Space	OS/2 boot files (start up); C: drive	SAVOBJ OBJ(SERVER11) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) or SAVLIB *ALLUSR or *NONSYS
QFPINT for Primary Language, QSYSnnnn for Secondary	QFPBSYS2	Server Storage Space	OS/2 Warp operating system (shared drive); D: drive	SAVOBJ OBJ(QFPBSYS2) LIB(QFPINT) DEV(tape-device-name) OBJTYPE(*SVRSTG) or SAVLIB *IBM or *NONSYS
QUSRSYS	SERVER13	Server Storage Space	OS/2 setup and log files; E: drive	SAVOBJ OBJ(SERVER13) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) or SAVLIB *ALLUSR or *NONSYS
QUSRSYS	QYNASYS1	Server Storage Space	Notes integration programs (shared drive); F: drive	SAVOBJ OBJ(QYNASYS1) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) or SAVLIB *ALLUSR or *NONSYS
QUSRSYS	QYNALNCD	Server Storage Space	Software images to be installed (shared drives); G: drive	SAVOBJ OBJ(QYNALNCD) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) or SAVLIB *ALLUSR or *NONSYS
/QFPNWSSTG	Various	Network Server Storage Space	Notes Server and User data	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QFPNWSSTG/SERVER100')
Note: Saving a network storage space saves the entire space as a single object. You need to vary off the Integrated PC Server (FSIOP).				

15.7.2 Backup and Recovery

You can divide the backup of Lotus Notes on the IPCS into two distinct parts:

- Part One

You can use standard AS/400 functions to back up and recover the configuration objects and storage spaces associated with the Integrated PC Server. These storage spaces include the network server storage space that contains the Notes server and Notes user databases. If you save the entire disk drive (which is the easiest and fastest approach to a disaster recovery plan), you cannot easily restore individual Notes databases or documents.

- Part Two

This part involves backing up the individual Notes databases, which is done by using the ADSTAR Distributed Storage Manager product (ADSM).

This section focuses mainly on Part One.

There are several AS/400 objects associated with Lotus Notes that need to be backed up using AS/400 functions. These objects include:

- The network server description and its associated lines, controllers, and devices
- The network server storage spaces associated with the network server description
- The storage spaces shared by all Integrated PC Servers on the AS/400 system (also known as the server storage spaces)
- The libraries or products associated with the Notes installation

15.7.3 Saving a Network Server Description

The network server description and its attached lines, controllers, and devices are backed up using these AS/400 save methods:

- Option 21 from the SAVE menu (also known as a full system save)
- The Save System (SAVSYS) command, as described in Section 5.3, "Omitting Objects on a SAVSYS Operation" on page 53
- The Save Configuration command, SAVCFG

None of these commands allow you to save a specific network server description. You can use the Retrieve Configuration Source (RTVCFGSRC) command to save the specific network server description. However, we do not recommend using RTVCFGSRC as it does not reflect the result of running the Install Network Server Application (INSNWSAPP) command. We recommend that you use any of the AS/400 methods described earlier.

Recommendation

To ensure a successful save of the network server description, vary off all network servers before you issue any save commands.

15.7.4 Saving the Server Storage Spaces

Each network server description has two server storage spaces (*SVRSTG objects) associated with it. These server storage spaces hold the entire disk volume that is created when you create the network server description. The server storage spaces are stored in the QUSRSYS library.

The names of the server storage spaces, which are the same as the names of the network server description, are followed by a "1" or a "3." The name of the C: drive storage space consists of the name of the network server description with a "1" added to it (for example, NTTEST1 for a storage space associated with a network server description named NTTEST). The name of the E: drive storage space consists of the name of the server description with a "3" added to it (for example, NTTEST3).

To save the server storage spaces, follow these steps:

1. Vary off the network server.
2. Enter the command:

```
SAVOBJ OBJ(NTTEST1 NTTEST3) LIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*SVRSTG)
```

15.7.5 Saving the Network Server Storage Spaces

When you set up the Integrated PC Server for the Notes server, you create a network server storage space for the Notes server and user data. You may have additional network server storage spaces on your AS/400 system. For example, if you have more than one Integrated PC Server, you have at least one network server storage space for each Integrated PC Server.

These network storage spaces contain user data. Save these spaces frequently since the user data is considered volatile.

The data in a network server storage space exists in the form of byte stream files. These objects exist in the root (/) file system.

Use the SAV command to save storage spaces in the integrated file system. Perform these steps:

1. Vary off the network server.
2. Enter the command:

```
SAV DEV('/QSYS.LIB/library-name.LIB/SAVF1.FILE')
OBJ('/QFPNWSSTG/space-name')
```

In this case, you are saving to a save file. For the OBJ parameter, specify the path name for each network storage space that you want to save. Network storage spaces are in the directory /QFPNWSSTG. Specify /QFPNWSSTG/xxxx, where xxxx is the name of the network storage space.

We do not describe the Save Licensed Program (SAVLICPGM) or the SAVLIB command in this section. They are described in the *Backup and Recovery* manual, SC41-5304, and in this redbook, Chapter 7, "Licensed Program and PRPQ Backup and Recovery" on page 85. There are no special requirements for Lotus Notes.

15.7.6 Restoring a Network Server Description

You can restore an individual network server description or a group of network server descriptions. To restore a network server description, enter:

```
RSTCFG OBJ(SERVER1)
DEV(tape-device-name) OBJTYPE(*NWS)
```

Note: In this example, SERVER1 is the network server description that was previously saved.

15.7.7 Restoring the Server Storage Spaces

You can restore the C: or the E: drive server storage space.

To restore the server storage space, use the RSTOBJ command as follows:

```
RSTOBJ OBJ(NTTEST1 NTTEST3)
SAVLIB(QUSRSYS) DEV(tape-device-name)
OBJTYPE(*SVRSTG)
```

In this example, the OBJ parameters, NTTEST1 and NTTEST3, name the server descriptions.

15.7.8 Restoring the Network Server Storage Spaces

If the Notes server files or user data become damaged, you can recover them by restoring a network storage space saved previously.

Use the integrated file system RST command to restore the network storage space as shown in this example:

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/QFPNWSSTG/space-name')
```

Recommendation

If you perform a complete save of the directories with the Integrated PC Server (FSIOP) varied off, your product is restored. Perform the following steps to complete the recovery of these products:

1. Add the links for the server descriptions by typing the following for each one:

```
ADDNWSSTGL NWSSTG(storage-name)
           NWS(network-server-description-name)
```

2. Vary on your Integrated PC Servers (FSIOP) by typing:

```
WRKCFGSTS *NWS
```

Select option 1 to vary on each Integrated PC Server (FSIOP).

Note: If you save the server storage space beneath QFPNWSSTG by using the command SAV OBJ('/QFPNWSSTG/server-storage'), the QFPNWSSTG must be created first. Create /QFPNWSSTG by completing these steps:

1. To create the server storage, enter the command:

```
CRTNWSSTG
```

2. Enter: RST OBJ('/QFPNWSSTG/server-storage)

3. To add the storage link, enter the ADDNWSSTGL command.

4. To vary on the Integrated PC Server (FSIOP), type:

```
WRKCFGSTS *NWS
```

Select option 1 to vary on.

15.7.9 Saving by Using the ADSM OS/2 Lotus Notes Agent

You can back up and recover a storage space as a whole unit using the AS/400 commands. However, you cannot back up and recover individual files in the storage space with these commands (such as a Notes database).

The ADSM Notes backup agent is an application that allows the backup and restore of individual documents within a Notes database. The ADSM Notes backup agent uses Lotus Notes APIs on the client to interface with Notes and ADSM client APIs, which interface with an ADSM server. It is *only available on an OS/2 Notes platform*.

The Notes backup agent supports:

- Performing an incremental backup of Notesxxx.NSF databases, and backing up only documents that changed since the previous incremental backup
- Restoring individual documents to a database that was previously backed up with the Notes backup agent

- Restoring individual documents that were deleted from a Notes database
- Restoring an entire Notes database
- Performing an incremental restore function by merging changed documents into a previous full database restore

In the following discussion, we use the ADSM Lotus Notes backup agent.

Note: Remember, the Lotus Notes backup agent can only save or restore Notes databases. If you do not have ADSM, you can also restore the K: drive to a different drive (for example, the L: drive) and issue the SBMNWSCMD CMD(XCOPY K: L: /s) command.

You must use the ADSM OS/2 Client to save or restore other PC files.

15.7.10 Backing Up Data Using the Notes Backup Agent

The Notes backup agent has two user interfaces:

- A command line interface
- A graphical workstation interface (GUI)

The command line interface operates from the AS/400 SBMNWSCMD command, the Notes Scheduler, or the Notes remote server console. You can back up and restore from the command line interface.

The workstation interface operates as an option in the Notes workspace on a workstation and is only used for restoring individual documents and databases.

You can perform an incremental backup of a NSF database by using the Notes backup agent. The Notes backup agent contains a program, called DSMNOTES, that can be specified on the SBMNWSCMD command.

The next example shows an incremental backup of a Notes database using the AS/400 command line interface with the Lotus Notes backup agent.

```
SBMNWSCMD  CMD('dsmnotes incr k:/notes/data/euorders.nsf-adsmpw=sckp')
            SERVER(LIMERICK) SVRTYPE(*BASE)
```

The DSMNOTES command can use a physical or logical path to find the database that you want to back up. For more information on DSMNOTES, refer to *OS/400 Integration of Lotus Notes V4R1*, SC41-5431.

15.7.11 Restoring Databases Using the Notes Agent

With the Notes agent, you have to use two different interfaces depending on what you are restoring:

- A complete database
- Merging a backup with the current version on disk

We use the Submit Network Server Command (SBMNWSCMD) to run the restore. Note that the ADSM command is DSMNOTES.

Here is an example running of the DSMNOTES command using SBMNWSCMD:

```
SBMNWSCMD  CMD('dsmnotes restore euorders.nsf -merge=yes -adsmpw=dmt')
            SERVER(LIMERICK) SVRTYPE(*BASE)
```

Note: In this example, merge is set equal to “yes,” which tells the Lotus backup agent to restore documents to an existing database.

15.7.12 Restoring Individual Documents

Use the Notes workspace interface if you want to restore:

- Individual documents
- Groups of documents
- Deleted documents

To restore individual documents to a Notes database, install the Notes agent on an OS/2 PC first. Remember that this Notes backup agent is for OS/2 clients. For more information on restoring these individual or group documents and deleted objects, see *OS/400 Integration of Lotus Notes V4R1*, SC41-5431.

15.8 NetWare on the Integrated PC Server

OS/400 Enhanced Integration for Novell NetWare provides NetWare services for AS/400 users that link them to either PC-based or Integrated PC Server-based NetWare servers. It also provides the interface to manage the QNetware file system.

This section describes how to back up and restore NetWare data on the Integrated PC Server.

15.8.1 QNetWare Characteristics

The QNetWare file system has several characteristics, which include:

- Providing access to data stored on a local or remote Integrated PC Server running Novell NetWare 4.10. It also offers access to data stored on stand-alone PC servers running Novell Netware 3.12 and 4.10.
- Providing access to NetWare Directory Services** (NDS**) objects.
- Providing the ability to dynamically mount NetWare file systems over any local mountable file system.
- Storing the data in stream files.

Note: The QNetWare file system is available only when NetWare Enhanced Integration for AS/400, OS/400 option 25, is installed on the system.

15.8.2 Network Server Storage Spaces and Volumes

The network server storage spaces can be linked to disk storage allocated for data, directories, and files. NetWare uses and manages this disk storage as a device on a PC. You can partition and create volumes on the network server storage space. To manage volumes from the AS/400 system, use the Work With NetWare Volumes (WRKNTWVOL) command, which requires Enhanced Integration for NetWare. You can also manage volumes using the NetWare INSTALL utility from RCONSOLE.

The NetWare system volume (SYS:) holds the server NetWare Loadable Module (NLMs), other NetWare NLMs, and user utilities. This volume is created using the Install Netware Server (INSNTWSVR) command. The system volume is unique to a NetWare server, and holds the NDS information and spool files.

15.8.3 Save and Restore Overview

The following table shows where the libraries and directories of the NetWare and OS/2 data are located.

<i>Table 32. Table Showing Libraries and the Objects for NetWare & OS/2</i>				
Types of Storage Spaces	Libraries and Directories	Drive Letters and Object Names	Object Content	Save Command
Server Storage Spaces	QUSRSYS	C: drive server1	OS/2 boot disk	SAVOBJ OBJ(NTTEST1) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) SAVLIB(*ALLUSR or *NONSYS)
Server Storage Spaces	QFPINT or QSYSnnnn	D: drive QFBSYS2	OS/2 disk	SAVLIB LIB(*IBM) or SAVLIB(*NONSYS)
Server Storage Spaces	QUSRSYS	E: drive server3	Netware boot drive	SAVOBJ OBJ(NTTEST3) LIB(QUSRSYS) DEV(tape-device-name) OBJTYPE(*SVRSTG) SAVLIB(*ALLUSR or *NONSYS)
Server Storage Spaces	QFPINT or QSYSnnnn	F: drive QFBSYS4	Netware programs	SAVLIB LIB(*IBM) or SAVLIB(*NONSYS)
Network Server Storage Spaces	/QFPNWSSTG	System	Netware system volume	SAV DEV('QSYS.LIB/tape-device-name.DEVD') OBJ('/QFPNWSSTG/NWRPL')
Network Server Storage Spaces	/QFPNWSSTG	User	Netware user volume	SAV DEV('QSYS.LIB/tape-device-name.DEVD') OBJ('/QFPNWSSTG/NWRPL')
AS/400	QSYS	NWSD	Server description	SAVCFG

15.8.4 Types of Storage Spaces

Server spaces contain your NetWare data. For each NWSD, you have at least two or more storage spaces. Normally, you can separate the spaces depending on whether the data is user data or more static, such as program data.

There are two types of storage spaces:

1. Server storage spaces

Each of these storage spaces is represented by a drive letter:

- C: drive

The C: drive is automatically created when you issue the Create NetWork Server Description (CRTNWSD) command. The Install NetWare Server (INSNTWSVR) command calls the CRTNWSD command.

This drive contains the data needed for booting OS/2, including CONFIG.SYS and the STARTUP.CMD files. The files on the C: drive are read-only format. The C: drive is named xxxxxx1, where xxxxxx is the name of the NWSD.

- D: drive

The D: drive holds the AS/400 copy of OS/2. It is created when installing the Integration Services for FSIOP feature. It holds the OS/2 kernel and other OS/2 programs used to start the Integrated PC Server. This drive is also in read-only format.

The AS/400 object name of the D: drive is QFPBSYS2. It is located in the QFPINT library. QFPINT is the library for the Integration Services for the FSIOP feature. For any other language, the D: drive is located in a library named QSYSnnnn, where nnnn is the language version number.

- E: drive

The E: drive contains the SERVER.EXE and STARTUP.NCF files that are used to start the server. It is created automatically when you issue the CRTNWSD command, which is also prompted by the INSNTWSVR

command. The E: drive is named *server3*, where “server” is the name of the NWSD. This storage space resides in the QUSRSYS library.

The E: drive is also used for some NetWare operations, such as installing licenses for a NetWare Loadable Module (NLM) or loading NetWare fixes. NetWare documentation refers to this drive as the local DOS drive.

- F: drive

This drive holds the AS/400 copy of NetWare and the NLMs that run NetWare on the Integrated PC Server.

The F: drive is also a read-only formatted drive. OS/400 applies PTFs to this storage space that fixes both Integrated PC Server NLMs and the network server programs.

The AS/400 name of this object is QFBSYS4, and it resides in the QFPINT library for the primary language. For non-US English languages, the library is QSYSnnnn, where nnnn is the language version number.

2. The network server storage spaces

The network server storage spaces contain the system volume and other volumes that you create to contain your directories and files, including AUTOEXEC.NCF. These storage spaces are stored in the integrated files system directory /QFPNWSSTG.

When you save data from this directory, you must save the entire storage space. Naturally, you cannot restore individual files and directories that were saved from this directory.

Note that the NWSD must be varied off to save /QFPNWSSTG.

15.8.5 Save and Restore Options

This section describes how to back up and restore NetWare data on both PC-based and Integrated PC Server-based NetWare servers. The following table illustrates the options for saving your NetWare data and implementing your save and restore strategy.

Command or Application	Server Storage Spaces				Network Server Storage Spaces		QSYS.LIB
	C: Drive	D: Drive	E: Drive	F: Drive	SYS Volume	User Volume	NWSD
SAVLIB *ALLUSR	X		X				
SAVLIB *IBM		X		X			
SAVCFG							X
SBACKUP or ARCserve					X	X	
SAV../QFPNWSSTG					X	X	
SAV../QNetWare					X	X	

Generally, we divide the save and restore process into two distinct options:

- Save everything
- Save specific objects associated with the NetWare components

We discuss each option in the following sections.

15.8.6 Saving Everything

This section describes two options that you can use to save everything. The option that you choose depends on what you want to save.

- Use option 21 from the SAVE Menu to save the entire system. This also saves the server and network server storage spaces and the Network Server Storage. The contents of the save media depend on whether a NWSD is varied on or off.

Saving with the NWSD varied *on* enables you to save information from the QNetWare directories. The advantage of this is that you can restore individual files.

Saving with the NWSD varied *off* enables you to save information within the /QFPNWSSTG directory. However, the individual files and directories cannot be restored from this save.

- Use option 23 to save user data. Under the covers of this menu option are these CL commands:

```
SAVSECDTA
SAVCFG
SAVLIB LIB(*ALLUSR) ACCPTH(*YES)
SAVDLO DLO(*ALL) FLR(*ANY)
SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ(('/*'
('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT)
('/QIBM/ProdData' *OMIT) ('/QOpenSys/QIBM/ProdData' *OMIT))
UPDHST(*YES)
```

This option provides a complete save scenario. Again, as in saving the entire system, varying off or on the NWSD during the save results in whether you are saving the /QFPNWSSTG or /QNetWare directories.

15.8.7 Saving Specific Objects

This section discusses saving the various components of NetWare, such as the NetWare configuration, server storage spaces, network storage spaces, NetWare volumes, and NetWare data.

15.8.7.1 Saving the NetWare Configuration Only

To save the Integration for NetWare configuration, you need to save the NWSD, as well as your line, IPX descriptions, and IPX circuit entries.

- To save the NWSD:
 1. Vary off the Integrated PC Server.
 2. Enter the command:

```
SAVCFG DEV(tape-device-name)
```
- Save the line, IPX descriptions, and IPX circuit entries

The following summarizes the steps to retrieve this information into a source file.

1. Create a source into which you can place the information. Enter:

```
CRTSRCPF FILE(MYLIB/NTTESTSRC)
```

2. Retrieve the configuration source of your NWSD. Enter:

```
RTVCFGSRC CFGD(BASELAN)
CFGTYPE(*NWSD) SRCFILE(MYLIB/NTTESTSRC)
SRCMBR(NWSRC) RTVOPT(*NET)
```

Note: In this example, BASELAN is the network server description name.

- Retrieve the configuration source from your IPX description into a source file. Enter:

```
RTVCFGSRG CFGD(BASELAN)
          CFGTYPE(*IPXD) SRCFILE(MYLIB/NTTESTSRC)
          SRCMBR(IPSRC) RTVOPT(*NET)
```

- Issue the command WRKIPXCCT, which provides the panel:

Work with IPX Circuits				
				System: SYSTEMXX
Type options, press Enter.				
1=Add 2=Change 4=Remove 5=Display 7=Display associated services				
8=Display associated routes 9=Start 10=End				
Opt	Circuit Name	Line Description	Line Type	Circuit Status
<u>5</u>	IPXCCT	ITSCTRN	*TRLAN	Inactive
—	NEWIPXCCT	NTTEST01	*TRLAN	Inactive

Figure 84. Work With IPX Circuits Display

- From this panel, choose option 5 to display the specific entries. Make note of the internal IPX number, frame type, and line to which they are connected. Refer to the following as an example of these displays.

Display IPX Circuit	
Circuit name	IPXCCT
Circuit status	Inactive
Network server description	BASELAN
Line description	ITSCTRN
Line type	*TRLAN
IPX Network number	429F2022
Frametype	*SSAP
Node address	400000002627
Enable for NLSP	*YES
MAC channel for NLSP	*BROADCAST
Router priority for NLSP	44
Cost override for NLSP	*CALC
Default maximum datagram size	*LIND
Throughput	*CALC
Delay	*CALC
Automatic start	*YES
RIP:	
State	*AUTO
Update interval	60
Age multiplier	4
SAP:	
State	*AUTO
Update interval	60
Age multiplier	4

Figure 85. Display IPX Circuit

6. Change the source files type from CL to CLP. Enter:

```
WRKMBRPDM FILE(MYLIB/XXXX)
```

This allows the files to be compiled and run.

7. Add the ADDIPXCCT statements into the IPSRC source file. This automatically creates the IPX circuit entry when the command is run.

```
Columns . . . : 1 71          Browse          hlaust/QCLSRC
SEU==> -
FMT ** - ...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7
***** Beginning of data *****
0001.00 /*          BASELAN 12/14/97 17:45:06 */
0002.00 CRTNWSN NWSN(BASELAN) RSRNAME(CC09) ONLINE(*YES) VRYWAIT(*NOWAIT) +
0003.00          LNGVER(2924) NTB(QNTBIBM) TEXT('IPCS - Base Lan Support') +
0004.00          CNTRYCODE(1) CODEPAGE(850) TYPE(*BASE) MSGQ(*JOBLOG) +
0005.00          CFGFILE(*NONE) STRNTB(*YES) STRTCP(*NO) TCPPORTCFG(*NONE) +
0006.00          TCPRTE(*NONE) TCPHOSTNAM(*NWSN) TCPDMNNAME(*SYS) +
0007.00          TCPNAMSVR(*SYS) SYNCTIME(*YES)
0008.00 CRTLINTRN LIND(ITSCTRN) RSRNAME(*NWSN) VRYWAIT(*NOWAIT) MAXCTL(40) +
0009.00          NWS(BASELAN 1) LINESPEED(16M) DUPLEX(*HALF) +
0010.00          MAXFRAME(1994) ACTLANMGR(*YES) TRNLOGLVL(*OFF) +
0011.00          TRNMGRMODE(*OBSERVING) LOGCFGCHG(*LOG) TRNINFBNC(*YES) +
0012.00          ADPTADR(400000002027) EXCHID(056FD260) +
0013.00          SSAP((04 *MAXFRAME *SNA)(12 *MAXFRAME *NONSNA)(AA +
0014.00          *MAXFRAME *NONSNA)(C8 *MAXFRAME *HPR)) ELYTKNRLS(*NO) +
0015.00          THRESHOLD(*OFF) LINKSPEED(16M) COSTCNN(0) COSTBYTE(0) +
0016.00          SECURITY(*NONSECURE) PRPDLY(*LAN) USRDFN1(128) +
0017.00          USRDFN2(128) USRDFN3(128) AUTOCRTCTL(*YES) +
0018.00          AUTODLTCTL(1440) CMNRCYLMT(2 5) +
0019.00          TEXT('Description for ITSC Token Ring Line')
0020.00 CRTCTLNET CTLD(ITSCTNET00) ONLINE(*NO) LINE(ITSCTRN) CNRSPTMR(170) +
0021.00          TEXT('CREATED BY AUTO-CONFIGURATION')
0022.00 CRTDEVNET DEVD(ITSCTIPX00) TYPE(*IPX) ONLINE(*NO) CTL(ITSCTNET00) +
0023.00          TEXT('CREATED BY AUTO-CONFIGURATION')
0024.00 CRTDEVNET DEVD(ITSCTTCP00) TYPE(*TCP/IP) ONLINE(*NO) CTL(ITSCTNET00) +
0025.00          TEXT('CREATED BY AUTO-CONFIGURATION')
0026.00 CRTDEVNET DEVD(ITSCTUSR00) TYPE(*USRDFN) ONLINE(*NO) +
0027.00          CTL(ITSCTNET00) TEXT('CREATED BY AUTO-CONFIGURATION')
0028.00 ADDIPXCCT CCTNAME(IPXCCT) LIND(ITSCTRN) IPXNETNBR(429F2001) +
0029.00          FRAMETYPE(*SSAP)
***** End of data *****

F3=Exit  F5=Refresh  F9=Retrieve  F10=Cursor  F11=Toggle  F12=Cancel
F16=Repeat find      F24=More keys
```

Figure 86. IPSRC Source File Display

8. Compile the two source members in the normal way. Call the compiled CL program NWOBJ and IPOBJ respectively.

Note: You can also use the SAVOBJ command to save the IPX circuit entries without retrieving the configuration source into a source file. To save IPX circuit entries, enter:

```
SAVOBJ OBJ(QAZSP*) LIB(QUSRSYS)
DEV(*SAVF) OBJTYPE(*FILE)
SAVF(MYLIB/IPSAVF)
```

The details of the restore operation are discussed in Section 15.8.11.1, “Restoring the NetWare Configuration” on page 282, of this redbook. Suffice to say, there are two methods of recovery. The first uses the INSNTWSVR, and the second uses the retrieve configuration method.

15.8.8 Saving the Server Storage Spaces

By saving the Server Storage Spaces, you save the:

- NetWare server configuration (C: and E: drives)
- Integration for NetWare Licensed programs (D: and F: drives)

These are discussed in the next two sections.

15.8.8.1 Saving the NetWare Server Configuration

You need the SAVOBJ command to save the NetWare server configuration. To save the configuration, complete these steps:

1. Vary off the Integrated PC Server.
2. Issue the SAVOBJ command to save the C: drive:

```
SAVOBJ OBJ(SERVER11) LIB(QUSRSYS)
      DEV(tape-device-name) OBJTYPE(*SVRSTG)
```

Note: In this example, the server storage space name is SERVER1. We tag a 1 to the end of this name for the C: drive. Therefore, the object is now SERVER11.

3. To save the E: drive, enter the command:

```
SAVOBJ OBJ(SERVER13) LIB(QUSRSYS)
      DEV(tape-device-name) OBJTYPE(*SVRSTG)
```

Note: In this example, the server storage space name is SERVER1. Since this is the E: drive, we tag a 3 to the end of the name. Therefore, the object is now SERVER13.

15.8.8.2 Saving the Integration for NetWare programs

The Integration for NetWare programs are in your NetWare server’s D: and F: drives. These exist in library QFPINT for the primary language. For any language other than the primary language, the library is QSYSnnnn, where nnnn is the language code number.

If you are saving licensed programs for recovery, issue the command:

```
SAVLIB LIB(*IBM) DEV(tape-device-name)
```

or

```
SAVLIB LIB(*NONSYS) DEV(tape-device-name)
```

If you are saving licensed programs for distribution, issue the SAVLICPGM command.

15.8.9 Saving the Network Storage Spaces

There are two main areas of saving network storage spaces. The first is the saving of NetWare volumes. The second is the saving of individual files, the entire directory, or SAV.RST resources in the QNetWare file system.

This is a good time to check out the concept of volumes. Refer to *OS/400 Integrating AS/400 with Novell NetWare*, SC41-5124, for a thorough discussion on this topic.

You can either save the /QFPNWSSTG storage space or the /QNetWare directory. Because you are saving these objects using the integrated file system structure, you need to use the SAV command.

15.8.9.1 Saving NetWare Volumes from the /QFPNWSSTG Directory

Use the SAV command to save the entire SYS and user volumes that are stored in the Integrated PC Server.

You must vary off the Integrated PC Server to save the storage spaces as a single object. You cannot save or restore any individual files or directories from the NetWare server on the Integrated PC Server.

You can save a volume by saving all storage spaces that the volume spans. Because volumes can span multiple storage spaces, be sure to save complete volumes.

Perform these steps to save complete volumes:

1. Put your system in a restricted state.

Note: This ensures that everything is saved from the /QFPNWSSTG directory.

2. Vary off the Integrated PC Server using WRKCFGSTS(*NWS).

The NWSD or Integrated PC Server must be varied off to save the /QFPNWSSTG directory. The Work With Configuration Status display appears as follows.

```
Work with Configuration Status                                SYSTEMXX
                                                            12/10/97 14:31:50
Position to . . . . . _____ Starting characters

Type options, press Enter.
 1=Vary on   2=Vary off   5=Work with job   8=Work with description
 9=Display mode status 13=Work with APPN status...

Opt  Description      Status      Type
---  -
  ---  BASELAN          ACTIVE      *NWS
  ---  ITSCTRN          ACTIVE      *LIN
  ---  SYSTEMXX         ACTIVE      *CTL
  ---  SYSTEMXX         ACTIVE      *DEV
  ---  SYSTEMYY         ACTIVE      *CTL
  ---  SYSTEMYY         ACTIVE      *DEV
  ---  SYSTEMZZ         ACTIVE      *CTL
  ---  SYSTEMZZ         ACTIVE      *DEV
  ---  ITSCTNET00       ACTIVE      *CTL
```

Figure 87. Work With Configuration Status Display

3. Enter the save command:

```
SAV DEV('QSYS.LIB/tape-device-name.DEVD')  
OBJ('/QFPNWSSTG/NWRPL')
```

In this example, you are backing up a user storage space called NWRPL. All user storage spaces are in the /QFPNWSSTG directory.

15.8.9.2 Saving NetWare Data from the QNetWare Directory

To better understand the /QNetWare directory saves, we need a picture of the structure as shown in the following figure.

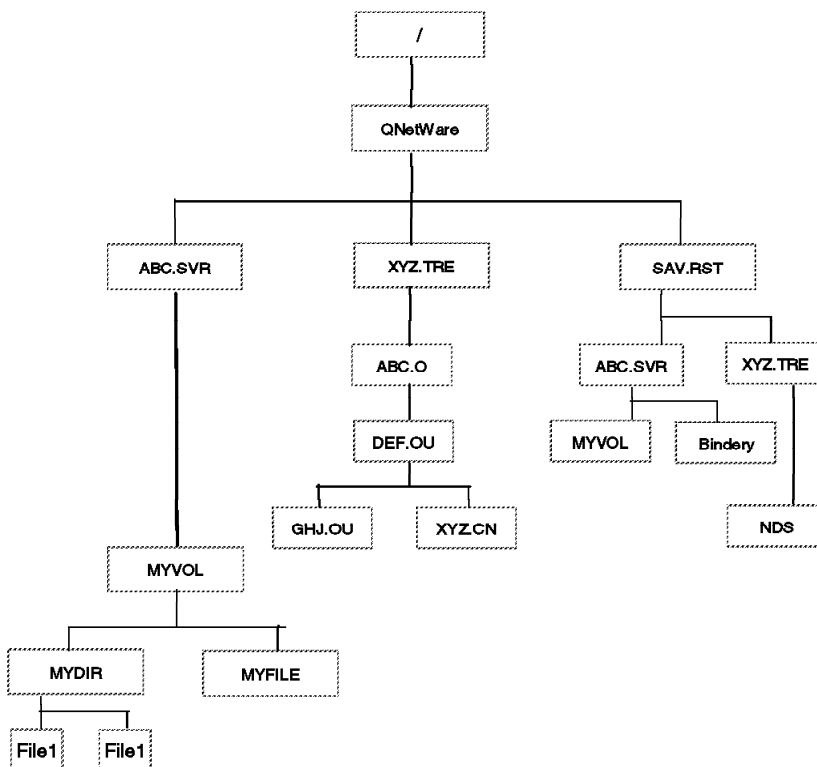


Figure 88. QNetWare File Structure

The previous structure represents the /QNetWare directory structure. It has multiple distinct file systems. This structure represents Novell NetWare servers and volumes in the format:

```
/QNetWare/SERVER.SVR/VOLUME
```

Note: When a volume under a server is accessed through the integrated file system menus, commands, or APIs, the root directory of the NetWare volume is automatically mounted on the VOLUME directory under /QNetWare. For a discussion on mounting, see Section 15.3.1, “Mounting User-Defined File System” on page 231.

QNetWare represents NDS trees on the network in the format:

```
/QNetWare/CORP_TREE.TRE/USA.C/ORG.O/ORG.UNIT.OU/SVR1_VOL.CN
```

The extension .TRE, .C, .O, .OU, and .CN represent NDS trees, countries, organizations, organizational units, and common names, respectively.

Note: If a Novell NetWare volume is accessed through the NDS path or an alias to a volume object, its root directory is also automatically mounted on the NDS object. For a discussion on mounting, see Section 15.3.1, “Mounting User-Defined File System” on page 231.

Use the integrated file system SAV command to save NetWare objects in the QNetWare directory. The objects you can save include: volumes, directories, files, or the entire QNetWare directory and resources in SAV.RST.

Follow these steps to save NetWare data through /QNetWare:

1. Create a save file to save NetWare objects. Enter:

```
CRTSAVF FILE(MYLIB/NETSAVE1)
```

2. Start an authenticated connection to the NetWare server. Enter:

```
STRNTWCNN SERVER(DOCSERV)  
SVRTYPE(*SERVER) CNNTYPE(*SAVRST)
```

Note: A *SAVRST connection requires that the following NLMs are present on the server, either already loaded or in the server’s search path:

```
NetWare 3.12 -- SMDR.NLM, TSA312.NLM  
NetWare 4.1 -- SMDR.NLM, TSANDS.NLM, TSA410(or TSA411).NLM
```

The *SAVRST connection must use the same user name and password used to start the *USER connection.

3. Make sure that you have a *SAVRST connection to the server before you save or restore NetWare objects.
4. To save the volume MYVOL on NetWare server ABC.SVR, enter:

```
SAV DEV(' QSYS.LIB/MYLIB.LIB/NETSAVE1.FILE')  
OBJ('/ QNETWARE/ABC.SVR/MYVOL')
```

5. To save specific directories or files in the volume, enter:

```
SAV DEV(' QSYS.LIB/MYLIB.LIB/NETSAVE1.FILE')  
OBJ('/ QNETWARE/ABC.SVR/MYVOL/MYDIR')
```

6. Resources include items such as an entire volume, the bindery on a NetWare 3.12 server, or the entire NDS for NetWare 4.1x. When you save resources under SAV.RST, the entire resource must be restored. To save NetWare resources such as “MYVOL,” enter:

```
SAV DEV(' QSYS.LIB/MYLIB.LIB/NETSAVF1.FILE')  
OBJ('/ QNETWARE/SAV.RST/ABC.SVR/MYVOL/')  
SYSTEM(*ALL) CLEAR(*ALL)
```

Note: Using SYSTEM(*ALL) ensures that the AS/400 system saves NetWare objects in both the local and remote servers. Using CLEAR(*ALL) ensures that the AS/400 system overrides any existing data in the save file during the save operation.

Note: Another option for backup is to use other backup alternatives such as SBACKUP or ARCserve for NetWare Version 6. The details are covered in *OS/400 Integrating AS/400 with Novell NetWare*, SC41-5124.

15.8.10 Restoring Everything

If you save using option 21 from the SAVE menu for a complete save of the system, use option 21 from the RESTORE menu to restore system and user data. These CL commands are initiated during an option 21 restore:

```
RSTUSRPRF
RSTCFG OBJ(*ALL)
RSTLIB SAVLIB(*NONSYS)
RSTDLO DLO(*ALL) SAVFLR(*ANY)
RST OBJ(('/*) ('/QSYS.LIB' *OMIT) ('/QDLS' *OMIT))
RSTAUT USRPRF(*ALL)
```

If you save using option 23, use the equivalent restore option to restore. This is option 23 from the RESTORE menu, which initiates the CL commands:

```
RSTUSRPRF USRPRF(*ALL)
RSTCFG OBJ(*ALL)
RSTLIB SAVLIB(*ALLUSR)
RSTDLO DLO(*ALL) SAVFLR(*ANY)
RST OBJ(('/*) ('/QSYS.LIB' *OMIT)
        ('/QDLS' *OMIT) ('/QIBM/ProdData' *OMIT)
        ('/QOpenSys/QIBM/ProdData' *OMIT))
RSTAUT USRPRF(*ALL)
```

15.8.11 Restoring Specific Objects

This section discusses restoring the various components of NetWare, such as the NetWare configuration, server storage spaces, network storage spaces, and NetWare data from the QNetWare directory.

15.8.11.1 Restoring the NetWare Configuration

Restoring the NetWare configuration consists of restoring both the network server description (NWSD), and your line, IPX descriptions, and IPX circuit entries.

- Restoring the NWSD

To restore the NWSD issue, enter the command:

```
RSTCFG OBJ(*ALL) DEV(tape-device-name)
OBJTYPE(*NWSD)
```

- Restoring the line, IPX descriptions, and IPX circuit entries

There are two methods to restore these configurations. The first is the retrieve configuration method, and the second is the INSNTWSVR method.

- Recovering using the RETRIEVE method:
 1. Ensure that you do *not* have a network server description, line description, IPX description, or IPX circuits with the server name.
 2. Run the CL program that you created during the save. Refer to Section 15.8.7.1, “Saving the NetWare Configuration Only” on page 275. Run the CALL NWOBJ program that you created in that section.
 3. Run the CL program that creates the network server description, line description, and links to the storage spaces. Run the CALL IPOBJ the program that you created in Section 15.8.7.1, “Saving the NetWare Configuration Only” on page 275.
 4. Make sure to include the ADDIPXCCT statements into the network source file.

- Recovering using the INSNTWSVR method:

1. Enter the command:

```
INSNTWSVR
```

This recreates your NetWare server configuration.

Note: Be sure that you do *not* have a network server description, line description, IPX description, or IPX circuits with the server name.

2. Specify VRYCFG(*NO) on the INSNTWSVR command. The Install NetWare Server (INSNTWSVR) display is shown as follows.

```

                                Install NetWare Server (INSNTWSVR)

Type choices, press Enter.

Network server description . . . _____ Name, *CPYONLY
Resource name . . . . . _____ Name
Port 1:
  Line type . . . . . _____ *ETH10M, *ETH100M, *TRN4M...
  Local adapter address . . . . _____ 020000000000-7FFFFFFFFF
  IPX circuit information:
  IPX network number . . . . . _____ 00000001-FFFFFFFF
  Frame type . . . . . _____ *SSAP *SSAP, *SNAP, *ETHV2, *ETHNTW
      + for more values _
Port 2:
  Line type . . . . . _____ *ETH10M, *ETH100M, *TRN4M...
  Local adapter address . . . . _____ 020000000000-7FFFFFFFFF
  IPX circuit information:
  IPX network number . . . . . _____ 00000001-FFFFFFFF
  Frame type . . . . . _____ *SSAP *SSAP, *SNAP, *ETHV2, *ETHNTW
Copy NetWare source . . . . . _____ *NO *YES, *NO
Vary on configuration . . . . . _____ *NO *YES, *NO
  IPX internal network number . . _____ *RANDOM 00000001-FFFFFFFE, *RANDOM
  Server message queue . . . . . _____ *JOBLOG Name, *JOBLOG, *NONE
  Library . . . . . _____ Name, *LIBL, *CURLIB
Network server storage space:
  Name . . . . . _____ *NWSD Name, *NWSD
  Size . . . . . _____ 200 100-8000
  FSIOP language version . . . . . _____ *PRIMARY *PRIMARY, 2963, 2966, 2980...
  Country code . . . . . _____ *LNGVER *LNGVER, 001, 002, 031, 033...
  Code page . . . . . _____ *LNGVER *LNGVER, 437, 850, 852, 857...

```

Figure 89. Install NetWare Server Display

3. Specify a replacement name for the network server storage area with a minimum size. Enter:

```
INSNTWSVR NWSD(DMTSD)
          CPYNTWSRC(*YES) NWSSTG(*NWSD 100)
```

4. Change the VRYCFG parameter in your NWSD back to *YES. Remove the link from the replacement storage space and make a link to the original spaces.

Note: To remove the link, use the RMVNWSSTGL command. To add a link, use ADDNWSSTGL command.

5. Be aware that if you install any special NLMs, name spaces on the original NetWare servers E: drive, or changes to STARTUP.NCF, they are not on the E: drive that you created with the INSNTWSVR command.

To restore your old E: drive, enter:

```
RSTOBJ OBJ(SERVER3) SAVLIB(QUSRSYS) DEV(tape-device-name)
```

Note: If you save the IPX circuit entries as described in Section 17.7.7.1, you can restore it using the command:

```
RSTOBJ OBJ(QAZSP*) SAVLIB(QUSRSYS)
DEV(*SAVF) OBJTYPE(*FILE)
SAVF(MYLIB/TEST1)
```

15.8.12 Restoring Server Storage Spaces

Restoring server storage spaces are grouped into two subsections. The first is restoring the NetWare server configuration, and the second is restoring the Integration for NetWare programs.

15.8.12.1 Restoring NetWare Server Configuration

NetWare server configurations are represented by the C: and E: drives. To restore them, perform these steps:

1. Vary off the Integrated PC Server.

2. To restore the C: drive, enter:

```
RSTOBJ OBJ(SERVER1) SAVLIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*SVRSTG)
```

3. To restore the E: drive, enter:

```
RSTOBJ OBJ(SERVER3) SAVLIB(QUSRSYS)
DEV(tape-device-name) OBJTYPE(*SVRSTG)
```

15.8.12.2 Restoring Integration for NetWare Programs

To restore integration for NetWare programs, use the RSTLIB command. Enter:

```
RSTLIB LIB(*IBM) DEV(tape-device-name)
```

or

```
RSTLIB LIB(*NONSYS) DEV(tape-device-name)
```

15.8.13 Restoring Network Storage Spaces

There are two distinct types of information in network storage spaces. The first is Netware volumes, which includes the Netware system and user volumes. The second includes the individual files, the entire directory, or the SAV.RST resources in the QNetWare file system.

15.8.13.1 Restoring NetWare Volumes from the /QFPNWSSTG Directory

You can save a volume by saving all storage spaces that the volume spans. Because volumes can span across multiple storage spaces, be sure complete volumes are saved and restored intact by saving and restoring all storage spaces associated with the volume. Use the Work With NetWare Volumes (WRKNTWVOL) command to see which storage spaces are spanned by a particular volume.

Perform these steps to restore the volumes:

1. Vary off the Integrated PC Server. Use the command:

```
WRKCFGSTS (*NWS) and take option 2
```

or

```
VRYCFG CFGOBJ(XXX) CFGTYPE(*NWS) STATUS(*OFF)
```

The variable XXX is the NWS name.

2. Enter the command:

```
RST DEV('/QSYS.LIB/MYLIB.LIB/NETSAVE1.FILE')  
OBJ('/QFPNWSSTG/NTEST1')
```

15.8.13.2 Restoring NetWare Data from the QNetWare Directory

The steps that you use to restore NetWare data from the QNetWare directory are equivalent to those that you took to save them. To restore, complete these steps:

1. Start an authenticated connection to the NetWare server.

```
STRNTWCNN SERVER(DOCSERV)  
SVRTYPE(*SERVER) CNNTYPE(*SAVRST)
```

2. Issue the RST command to restore volume MYVOL on NetWare server ABC.

```
RST DEV('/QSYS.LIB/MYLIB.LIB/NETSAVE1.FILE')  
OBJ('/QNETWARE/ABC.SVR/MYVOL')
```

3. For remote Integrated PC Server-based server, you need to specify an additional parameter on the RST command—the *SYSTEM(*ALL)* parameter:

```
RST DEV('/QSYS.LIB/MYLIB.LIB/NETSAVE1.FILE')  
OBJ('/QNETWARE/ABC.SVR/MYVOL')  
SYSTEM(*ALL)
```

15.8.13.3 Restoring NetWare Data from the /QNetWare Directory

You can restore individual objects in the QNetWare directory such as volumes, directories, or files. Or you can save the entire QNetWare directory. You can also restore resources within the SAV.RST directory.

Note: Another restore option is to use other backup alternatives such as SBACKUP or ARCserve for NetWare Version 6. The details are covered in *OS/400 Integrating AS/400 with Novell NetWare*, SC41-5124. SBACKUP is a NetWare backup utility and is used to perform incremental backups.

15.8.14 Other Tips and Techniques

To perform a backup of entire volumes of data for disaster recovery, you can use OS/400 commands like SAV. In this situation, we recommend that you do not store user data in SYS:. Instead, store it in a separate volume in a separate storage space. Separate storage spaces allow recovery of user data volumes without changing NDS, which is stored on SYS:. If NDS needs to be recovered, and there is not a copy of NDS on the network from which to recover, NDS can be recovered by restoring SYS:. Refer to Section 15.8.9, "Saving the Network Storage Spaces" on page 279, for details on saving and restore storage spaces.

To perform an incremental backup of data to a tape device, use a NetWare backup utility, such as SBACKUP. The tape device is attached to a PC-based server. Incremental backup of data to AS/400 media is not currently supported by OS/400 Integration for Novell NetWare.

To save storage spaces from /QFPNWSSTG directory, you need to vary off the Integrated PC Server. If the Integrated PC Server is varied on during the save, the storage spaces get locked. The Integrated PC Server may have changed

data cached in memory. The Licensed Internal Code involved in save operations knows about data changed in AS/400 memory, but not about data changed in memory on an Integrated PC Server. Even if the lock is not there, the resulting saved storage space may, or may not, be usable if it is restored.

Attention

If you perform your complete save of the directories with the Integrated PC Server (FSIOP) varied off, your Windows NT is restored. You need to perform the following steps to complete the recovery of these products:

1. To add the links for the server descriptions, type the following for each server description:

```
ADDNWSSTGL NWSSTG(storage-name)
           NWS(network-server-description-name)
```

2. To vary on your Integrated PC Servers (FSIOP), type:

```
WRKCFGSTS *NWS
```

Select option 1 to vary on each Integrated PC Server (FSIOP).

Note: If you save the server storage space beneath QFPNWSSTG by using the command SAV OBJ('/QFPNWSSTG/server-storage), the QFPNWSSTG must be created first. To create QFPNWSSTG, perform these steps:

1. To create the server storage, use the CRTNWSSTG command.
2. Enter: RST OBJ('/QFPNWSSTG/server-storage)
3. Add the storage link by using the ADDNWSSTGL command.
4. To vary on the Integrated PC Server (FSIOP), type: WRKCFGSTS *NWS. Select option 1 to vary on.

Another useful reminder is to make sure that you specify the SYSTEM(*ALL) or the SYSTEM(*RMT) parameter in the SAV command for saving data on the remote system.

We recommend that you put your system in a restricted state when saving either the /QFPNWSSTG directory or /QNetWare directory. It does not vary off the NetWare server which is still usable. However, since the NetWare server is active, you cannot restore the storage spaces associated with the Integrated PC Server. This includes the SYS: and any other volumes.

15.8.14.1 Saving and Restoring Internetwork Packet Exchange (IPX) Support

The AS/400 system ships IPX as part of OS/400. IPX allows the AS/400 system to act as an IPX router. Save and restore all IPX configuration files as a group. Many logical files are defined. Their names are given in the following table.

Table 34 (Page 1 of 2). IPX Logical Files

IPX Logical Files	Description
QAZSPLADR QUSRSYS	System-supplied IP over IPX address file that is defined on the QAZSPPADR physical file
QAZSPLCCT QUSRSYS	System-supplied IPX circuit logical file that is defined on the QAZSPPCCT physical file

<i>Table 34 (Page 2 of 2). IPX Logical Files</i>	
IPX Logical Files	Description
QAZSPLRTE QUSRSYS	System-supplied IPX route logical file that is defined on the QAZSPPRTE physical file
QAZSPLSRV QUSRSYS	System-supplied IPX service logical file that is defined on the QAZSPPSRV physical file

Saving just the physical files and restoring them causes problems. The logical files used by different functions to access the data stored in the physical files point to a renamed physical file if only a physical file is restored. The restore database functions create the renamed physical file to maintain the indexes to the logical files. Another reason to save and restore the configuration files as a group is that there are some dependencies between some of the files. Saving and restoring only a subset of the files can cause problems, especially when activating IPX processing.

To save all the IPX files, enter:

```
SAVOBJ OBJ(QAZSP*) LIB(QUSRSYS)
      DEV(tape-device-name) OBJTYPE(*FILE)
```

The equivalent restore command of these IPX files is:

```
RSTOBJ OBJ(QAZSP*) LIB(QUSRSYS)
      DEV(tape-device-name) OBJTYPE(*FILE)
```

The resources under the SAV.RST directory include binderies for NetWare 3.12 servers, NDS information for NetWare 4.1 trees, and volume for all servers. In regard to a bindery, this is the only place it shows up under /QNetWare. By saving a volume /QNetWare/SAV.RST/myserver.SVR/myvol as opposed to /QNetWare.myserver.svr/myvol, the entire volume is saved as a single object on the AS/400 system. This significantly improves performance, but does not allow individual files and directories to be restored. You can only restore the entire volume.

15.9 OS/2 Warp Server for AS/400 (Formerly Known as LAN Server for OS/400)

Note: OS/2 Warp Server for AS/400 was known as LAN Server for OS/400 prior to V4R1.

IBM OS/2 Warp Server for AS/400 is a client/server product that uses the Integrated PC Server (the FSIOP). It also works with PC software called LAN Requesters to provide fast file serving and print serving to your AS/400 system. In this regard, the IPCS is a piece of hardware, an I/O processor that you attach to your AS/400 system. It was originally known as the File Server IO processor (FSIOP). OS/2 Warp Server for AS/400 provides the interface to manage the /QLanSrv file system when the Licensed Program Product is installed.

Information stored by LAN Server for OS/400 appears in several different directories on the system.

- /QLanSrv

This directory is a logical representation of access to the information.

- /QFPNWSSTG

This is the physical storage of LAN Server for OS/400 information. In the /QFPNWSSTG directory, each logical PC drive appears as one large object to the AS/400 system. The virtual drives are the network server storage spaces. Use the directories with the /QFPNWSSTG to save entire virtual drives. Individual files and directories cannot be restored from saved copies of the directories within the /QFPNWSSTG directory.

- /QSYS.LIB/QXZ1.LIB

This directory holds the licensed program information that does not change. It is saved using SAVLIB (*IBM) or SAVLIB (*NONSYS). The /QSYS.LIB/QXZ1.LIB directory holds three objects:

```
QFPHSYS1.SVRSTG
QFPHSYS2.SVRSTG
QFPHSYS3.SVRSTG
```

- /QSYS.LIB/QUSRSYS.LIB

This directory holds copies of server storage spaces. QUSRSYS library is saved when you specify *ALLUSR or *NONSYS for the SAVLIB command.

- Library QSYS29nn (29nn is a language number)

This holds the secondary language licensed system code. This library holds the national language version of the code stored in QXZ1. It contains two objects:

```
QFPHSYS2.SVRSTG
QFPHSYS3.SVRSTG
```

Note: QFPHSYS1.SVRSTG does not have an NLS version.

- QLANSRVSR

This directory is a temporary storage area for files that are in the process of being saved to backup media.

15.9.1 OS/2 Warp Server for AS/400 Structure

The following figure describes the layout of the QLANSRV file structure.

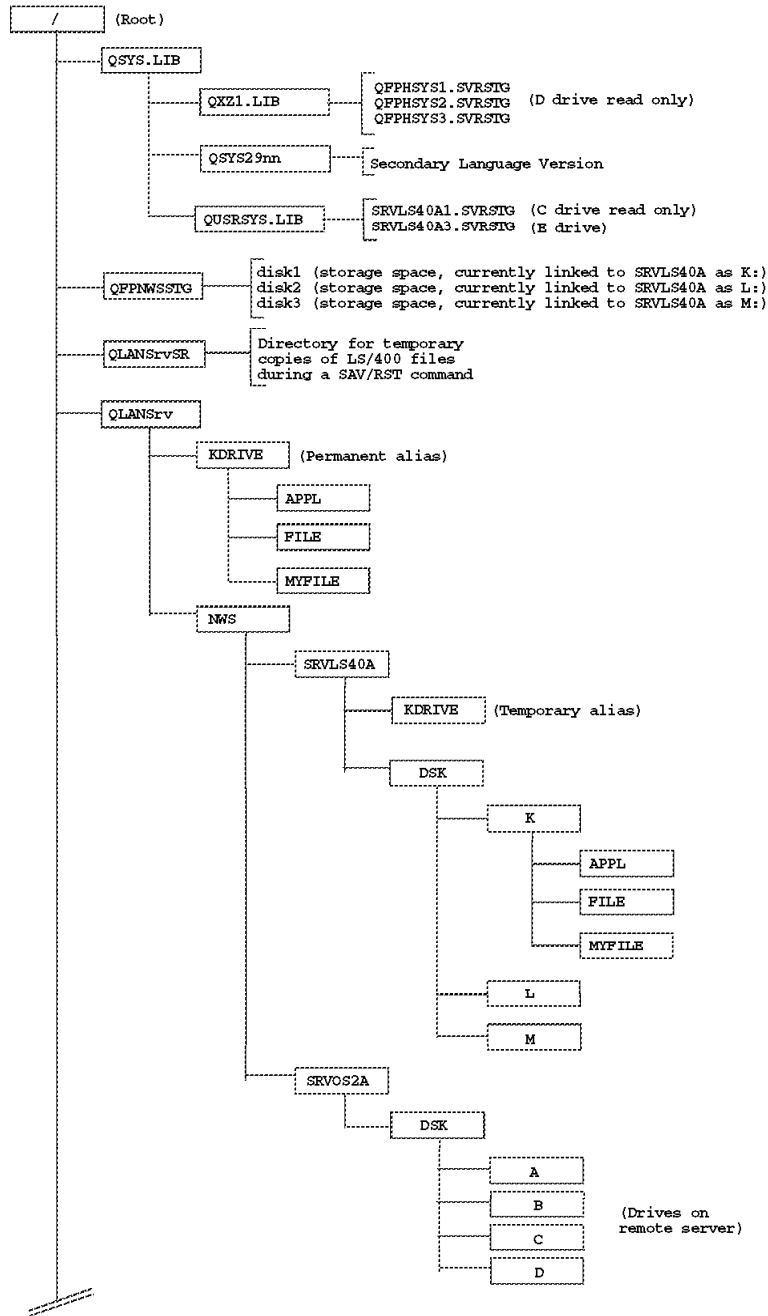


Figure 90. LAN Server/400 Integrated PC Server File System Objects

Note: The name change from LAN Server for OS/400 to OS/2 Warp Server for AS/400 occurred in V4R1. The product ID 576n-XZ1 remained the same on V4R2 as on V4R1 since the product was not refreshed to a V4R2 level. LAN Server for OS/400 was discontinued in V4R1.

The LAN Server for OS/400 product does not run on the IPCS-3, which was introduced in V3R7. It only runs on the IPCS-1. OS/2 Warp Server for AS/400 runs on all IOP hardware (IPCS and FSIOAs).

Note: The IPCS-1 is a 6506. The IPCS-3 is a 6616, which was introduced in V3R7. OS/2 Warp Server does not run on releases prior to V4R1.

15.9.2 Backup and Recovery for OS/2 Warp Server for AS/400

The save and restore of the OS/2 Warp Server for AS/400 file system data is done with integrated file system commands. The integrated file system commands SAV and RST, allow you to specify an object name using the integrated file system object naming convention. For more information about integrated file system commands, see Section 15.1, “The Integrated File System” on page 221.

The integrated file system design allows you to save files from both the local IPCS and the remote servers, whether they are IPCS or OS/2 Warp Server for AS/400 systems. This flexibility makes the AS/400 system a powerful component in the domain because you can use it to save the server data from any LAN server system in the domain.

15.9.3 Storage Spaces

When we discuss backup and restore procedures, we distinguish between two kinds of storage spaces: server and network.

- Server storage space

These storage spaces hold an entire File Access Translator (FAT) disk volume that is created when you create a network server description and are used by OS/2 Warp Server for AS/400. Server storage spaces contain licensed programs and system files, such as OS/2 code, LAN Server code, IPCS device drivers, IPCS administration applications, CONFIG.SYS, NET.ACC, SWAPPER.DAT, and dump files.

We use OS/400 SAVOBJ/RSTOBJ or SAVLIB/RSTLIB commands to perform save and restore for server storage spaces. This is because storage spaces are stored in libraries and not in the integrated file system structure.

- Network server storage space

These storage spaces are usually created and used by the OS/2 Warp Server for AS/400 administrator. They also can be created by users. Network server storage spaces hold the directories and files that make up the entire HPFS disk volume. In this context, the directories within the /QFPNWSSTG are called the network storage spaces.

Note: Throughout this section the term “storage space” refers to network server storage spaces, unless otherwise indicated.

15.9.4 Authority Requirements for Saves

The special authorities required to save the various OS/2 Warp Server for AS/400 components are:

- SAVSYS authority

This is required when saving using the /QFPNWSSTG directory.

- ALLOBJ authority

This is required when saving /QLanSrv objects and their authority information if both of the following conditions are met:

- You are a defined user in the LAN domain.
- The domain controller is a FSIOF (IPCS) on the local AS/400 system.

Authority information for OS/2 Warp Server for AS/400 objects is stored within the objects, not with the user profiles that have authority. The SAVSECDTA and SAVSYS commands do not save authority information for OS/2 Warp Server for AS/400 objects. The authority information is saved when you save the objects, if you have sufficient authority. If not, the object is saved without the authority information.

- LAN administrator authority

To save objects from the /QLanSrv directory, you need LAN administrator authority. This ensures that the file access control information is saved. To do this, use the command:

```
SBMNWSCMD CMD(' NET USER username /PRIVILEGE:ADMIN' )  
SERVER(network-server-description-name)
```

15.9.5 Restricted State

Before you can save and restore the local files for OS/2 Warp Server for AS/400, we recommend that you put the OS/2 Warp Server for AS/400 system in a restricted state. This ends all the subsystems running on the AS/400 system and leaves only the system console operational. If you cannot put the system in a restricted state, verify that no files are open by using the Work With NWS Sessions (WRKNWSSSN) command.

Put the AS/400 system in restricted state only when you want to save files that are stored on the AS/400 system itself. If you want to save files from a remote server, take the AS/400 system out of restricted state by restarting the subsystems. When the OS/2 Warp Server for AS/400 is in a restricted state, you **cannot** save remote data because you cannot start any of the system functions required to access the remote systems.

To save IPCS data in a restricted state, your AS/400 system must have either a domain controller or a backup domain controller configured. It is not possible to properly vary on a NWSD without access to a controller. If you have multiple Integrated PC Servers on your AS/400 system, only one of them must be a controller; the others can access the controller using the interconnect function.

15.9.6 Tips for Saving OS/2 Warp Server for AS/400 on RISC Machines

You can use option 21 or 23 from the SAVE menu or Option 11 from the BACKUP menu to save the OS/2 Warp Server for AS/400 environment. The system attempts to save the directories within the QLANSRV directory or the directories within the QFPNWSSTG directory. Note the following:

- If the network server description is varied on, the save contains /QLanSrv directories associated with the Integrated PC Server, also known as the FSIOP.
- If the network server description is varied off, the save contains directories within the /QFPNWSSTG directory. Individual files cannot be restored from a saved copy of a directory within the /QFPNWSSTG directory.
- If a faster save is required, you need to vary the Integrated PC Server off. From a performance standpoint, saving with the IPCS varied off is dramatically faster. You cannot restore individual files or directories from this save. Conversely a "slow" save results from leaving the Integrated PC Server varied on.

- We recommend that you put objects that change frequently (such as files) in one or two subdirectories in /QLanSrv. Save these directories frequently with the Integrated PC Server varied on.
- In addition, consider partitioning your volatile data (data that changes frequently) and static data (such as program data that changes less frequently) on separate network drives. For static data, save with the Integrated PC Server varied off. For volatile data, if you need to restore at the file or directory level, you need to save with the Integrated PC Server varied on. If you need to restore at the file or directory level, but find that varying on the Integrated PC Server does not provide the required performance, consider these steps:
 1. Save with the Integrated PC Server varied off.
 2. When a restore is required, create a temporary network drive in /QLanSrv, and restore into the temporary network drive.
 3. Selectively restore the required files from the temporary network drive.
 4. Delete the temporary drive when complete.
- When you save a directory within the /QFPNWSSTG directory, specify SUBTREE(*ALL), which is the default. These directories contain files that must be saved and restored as a group.
- Starting with V3R7 for RISC machines and V3R2 for CISC machines, there is an important change in the way the CHGPERIOD(*LASTSAVE) parameter on the SAV command works. When specifying the *LASTSAVE value, objects are saved that changed since the last time they were saved with the SAV command specified as UPDHST(*YES). It is important to note that prior to V3R7 and V3R2, using OS/2 Warp Server for AS/400 and saving with CHGPERIOD(*LASTSAVE) always completely saved OS/2 Warp Server for AS/400. This means *all* objects were saved for OS/2 Warp Server for AS/400, *not just the changed* objects.

15.9.7 Examples of Saving Specific OS/2 Warp Server for AS/400 Files

This section provides examples of saving objects within the OS/2 Warp Server for AS/400, including objects with multiple names, the domain controller database, specific objects, a directory, and network storage spaces.

15.9.7.1 Saving OS/2 Warp Server for AS/400 (LAN Server/400) Objects with Multiple Names

When the OS/2 Warp Server for AS/400 server objects have multiple names, the additional names are called aliases and netnames. Netnames are temporary and are defined during a session. Definitions for aliases are stored in the OS/2 Warp Server for AS/400 domain controller database (DCDB). This is similar to symbolic links in the QOpenSys file system.

When varying on the first network server description in the domain, the OS/2 Warp Server program creates directories for each of the aliases that are defined. When you vary on a network server description or a remote OS/2 Warp Server, the OS/2 Warp Server program creates directories for each of the netnames that is currently defined.

Objects are marked to ensure that you save the contents of an object only once, even if the object has only one name. If you save the entire /QLanSrv directory, you are saving each file and directory once even if you have aliases. To save

the nicknames (aliases) that have been set up on your system, you must save the DCDB. See Section 15.9.7.2, “Saving the Domain Controller Database (DCDB)” on page 293.

15.9.7.2 Saving the Domain Controller Database (DCDB)

If one of the network servers on your AS/400 system is the domain controller, you need to save the DCDB. There are two methods you can use:

- The DCDB replicator service to replicate the DCDB to a backup domain controller
- Save the server storage space that contains the DCDB directories
`SAVOBJ OBJ(QUSRSYS/server3) OBJTYPE(*SRVSTG)`

Note: For server3, substitute the name of the network server description that is followed by a “3.”

15.9.7.3 Saving Specific OS/2 Warp Server Objects

These specific examples are best summarized in the following table.

<i>Table 35. Examples of Saving Directory and Files</i>	
Commands	Description
<code>OBJ('/QLANSRv/NWS/FS11/DSK/T/FILES')</code>	Saving a specific file in a specific directory on a local AS/400
<code>OBJ('/QLANSRv/NWS/SERVER1/DSK/T/FILES/FILEA.TXT') SYSTEM(*RMT)</code>	Saving a specified file on a remote system

15.9.7.4 Saving a Directory for an Integrated PC Server

Saving an entire directory is similar to saving a library. You need to put your system in a restricted state to save IPCS directories. This ensures that everything is properly saved and no changes are made to the directory or its contents. However, the network server description must remain varied on during the save.

The following table summarizes the save actions.

<i>Table 36. Commands to Save All Local IPCS Directories</i>	
Commands	Description
<code>OBJ('/QLANSRv/*')</code>	Save all local IPCS directories
<code>OBJ('/QLANSRv/NWS/iop-name')</code>	Save a specific directory for a specific Integrated PC Server

Another way of viewing this is to use the AS/400 Operations Navigator, as shown in Figure 91 on page 294.

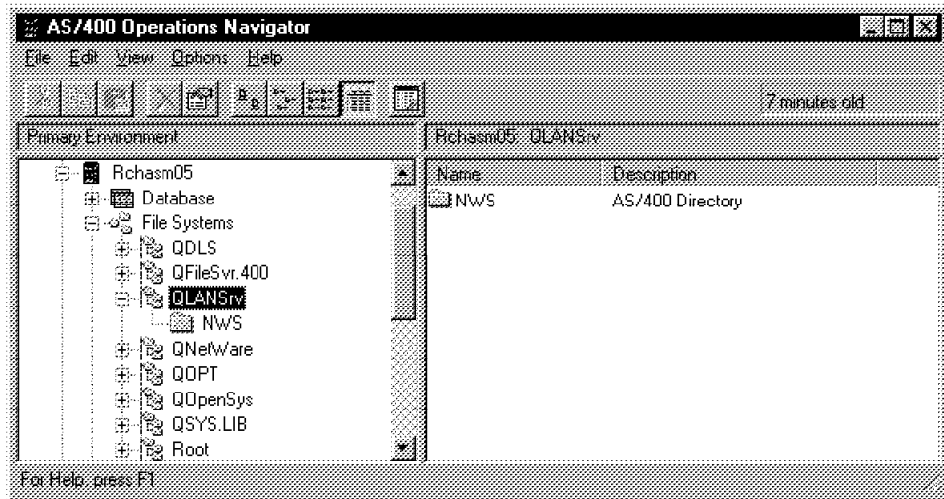


Figure 91. /QLanSrv Directory Structure Viewed from Operations Navigator

15.9.7.5 Saving Network Storage Spaces

You can save a storage space (/QFPNWSSTG) or move it to another system. You can also save it for faster recovery. See Section 15.9.6, “Tips for Saving OS/2 Warp Server for AS/400 on RISC Machines” on page 291, for more information.

When saving in this manner, vary off the IPCS. The table below summarizes the commands to save storage spaces.

Command	Description
OBJ('/QFPNWSSTG/drive-name')	To save a specific storage space
OBJ('/QFPNWSSTG/*')	To save all storage spaces

Figure 92 on page 295 shows the /QFPNWSSTG storage space as seen by the Operations Navigator.

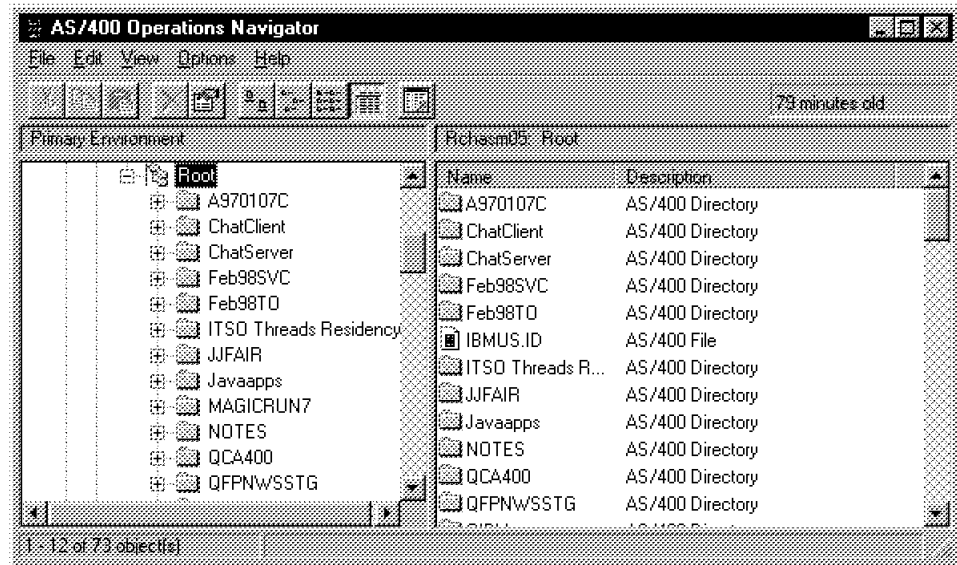


Figure 92. /QFPNWSSTG Directory Structure Viewed From Operations Navigator

15.9.8 Restoring OS/2 Warp Server for AS/400

Restoring the OS/2 Warp Server for AS/400 environment can be done with the IPCS varied on or off, as discussed in this section.

15.9.8.1 Recovery with the IPCS Varied Off

These steps are performed after running Restore Authority (RSTAUT) and applying PTFs that were not included in the last SAVSYS.

1. Add links for the server description using the Add Server Storage Link (ADDNWSSTGL) command. Enter:

```
ADDNWSSTGL NWSSTG(network-storage-name)
NWS(network-server-description-name)
```

2. Vary on the IPCS. Enter:

```
WRKCFGSTS *NWS
```

Select option 1 to vary on each IPCS.

15.9.8.2 Recovery for SAV Performed with the IPCS Varied On

Complete the following steps after performing a normal IPL at the end of restoring your system:

1. If the IPCS is varied on, vary it off.
2. Create any needed network storages by using the Create NWS Storage Space (CRTNWSSTG) command. Enter:

```
CRTNWSSTG NWSSTG(ABCD)
```

3. Add the storage links by using the Add Server Storage Link (ADDNWSSTGL) command. Enter:

```
ADDNWSSTGL
```

4. Vary on the IPCS using the command:

```
WRKCFGSTS *NWS
```

Select option 1.

5. Restore the LAN Server/400 data using the Restore command (RST).

```
RST('/QLANSrv')
```

Press Enter.

15.9.8.3 BACKACC and RESTACC

OS/2 Warp Server for AS/400 has two commands to save the access control information for the files stored on the server. In the OS/2 Warp Server for AS/400 environment, it is not necessary to use these commands if you make frequent backups of the E: drive.

However, under certain circumstances, you may want to back up the access control information separately. For example, you may not want to grant *ALLOBJ or administrator authority to your users, or you may have an OS/2 Warp Server for AS/400 as the domain controller. Either of these situations justify using the BACKACC command. We suggest that you use the BACKACC and RESTACC commands with the Submit Network Server Command (SBMNWSCMD). For example, enter:

```
SBMNWSCMD CMD('BACKACC K: /S')
```

This command backs up the access control information associated with the K: drive and the subdirectory below the K: drive in a default file called ACLBAKK.ACL. The last character of the filename before the period (in this case, K:) tells you which drive letter was backed up in this file. You need to execute this command for each drive that you want to back up.

Use the RST command if you need to restore the OS/2 Warp Server for AS/400 file system objects. Then use the RESTACC command to restore the file access control information.

Errors

When using the BACKACC command, there are two common sets of reported errors:

1. NET3550: File NETACC.OLD exists; NETACC.BKP may be damaged
NET3566: BACKACC cannot back up NET.ACC and NET.AUD NET3567:
BACKACC cannot backup the access control list information
2. NET3551: File NETAUD.OLD exists; NETAUD.BKP may be damaged
NET3566: BACKACC cannot backup NET.ACC and NET.AUD NET3567:
BACKACC cannot backup the access control list information

To resolve the problem, delete the file NETACC.OLD from the E:IBMLANACCOUNTS directory and delete the NETAUD.OLD file from the E:IBMLANLOGS directory of the server. Then, try the command again.

You can also use the following commands on the AS/400 system to delete these files (Note that you must enter each option as one command line.):

```
RMVLNK OBJLNK('/QLanSrv/NWS/network-server-description-name/IBMLAN$  
/ACCOUNTS/NETACC.OLD')
```

or

```
RMVLNK OBJLNK('/QLanSrv/NWS/network-server-description-name/IBMLAN$  
/LOGS/NETAUD.OLD')
```

The following table summarizes various save and restore options for the OS/2 Warp Server for AS/400 product.

<i>Table 38. Summary of Save and Restore Options</i>				
Objects Saved	Save Command	Status IPCS Varied On or Off	Restricted State AS/400 Yes or No	Remote Save Yes or No
Storage spaces located in /QFPNWSSTG and libraries QUSRSYS, QXZ1, and any national language version of QXZ1 used for disaster recovery for an entire system.	Option 21 from the SAVE menu	Off	Yes	No
Storage spaces of a specific network server located in /QFPNWSSTG on the local AS/400 system.	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QFPNWSSTG/DISK1')	Off	N/A	No
Files and Directories located in /QLanSrv used for disaster recovery and file restoration.	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/**') ('/QSYS.LIB' *OMIT) ('/QDLS *OMIT)) or SAVE menu Option 21	On	Yes	No
Files and directories located in /QLanSrv, changed or created within a date range. Used as incremental backup.	SAV DEV('/QSYS.LIB) OBJ('/QLanSrv/**') CHGPERIOD(mm/dd/yy)	On	Yes	No
Files and directories located in /QLanSrv of a remote system, changed or created within a date range. Used as incremental backup.	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QLanSrv/**') ('/QSYS.LIB' *OMIT) ('/QDLS *OMIT)) CHGPERIOD(mm/dd/yy) SYSTEM(*RMT)	On	No	Yes
Specific directories and files on a local system.	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QLanSrv/NWS/network-server-description-name/DSK/K/file') ('/QSYS.LIB' *OMIT) ('/QDLS *OMIT))	On	Yes	No
Specific directories and files on a remote system.	SAV DEV('/QSYS.LIB/tape-device-name.DEVD') OBJ('/QLanSrv/NWS/SRVOS2A/DSK/D/files') SYSTEM(*RMT) ('/QSYS.LIB' *OMIT) ('/QDLS *OMIT))	On	No	Yes
OS/2 Warp Server product library	SAVLIB *NONSYS or SAVLIB *IBM	N/A	Yes	N/A
Storage space containing the DCDB. The name of the object is the name of the server followed by a three.	SAVOBJ OBJ(network-server-description-name3) LIB(QUSRSYS) OBJTYPE(*SVRSTG)	Off	N/A	N/A

15.9.9 PC Client

With the save and restore functions of OS/2 Warp Server for AS/400 discussed in Section 17.8, you can manage the backup and recovery of data stored on LAN Server systems.

Managing the backup and recovery of other workstations in your organization requires other techniques for which OS/2 Warp Server for AS/400 was not designed. We recommend that you use another product called ADSTAR* Distributed Storage Manager (ADSM) to help you manage your workstations. See Section 9.1, "ADSTAR Distributed Storage Manager/400 (ADSM/400)" on page 109, for more information.

15.10 Firewall

The firewall is made up of several components, all of which must be saved. Backup and recovery is important even for your firewall. It provides a means to restore a configuration in the event of system failure or other catastrophic event. Integrate the strategy used with the firewall into your existing backup and recovery and disaster recovery plan.

There are many objects, which taken together, make up the firewall. These objects are stored in AS/400 libraries, integrated file system directories, and network server storage spaces.

Table 39. Firewall Objects—Description and Location

Object Name	Library or Integrated File System Directory	Object Type	Description
firewall	QSYS	*NWSD	Network server description for the firewall
firewall00	QSYS	*LIND	Line description for IPCS *INTERNAL port
firewall01	QSYS	*LIND	Line description for IPCS port 1
firewall02	QSYS	*LIND	Line description for IPCS port 2
firewNet	QSYS	*CTLD	Controller used by the AS/400 TCP/IP to communicate with the firewall through the *INTERNAL port
firewTCP	QSYS	*DEVD	Device used by the AS/400 TCP/IP to communicate with the firewall through the *INTERNAL port
firewall1	QUSRSYS	*SVRSTG	IPCS C: drive—OS/2 boot disk
QFPBSYS2	QFPINT	*SVRSTG	IPCS D: drive—OS/2 disk
firewall3	QUSRSYS	*SVRSTG	IPCS E: drive—TCP/IP configuration and firewall base configuration
QISASTG1	QIPSINT	*SVRSTG	IPCS F: drive—firewall programs
firewall00	Integrated File System Directory /QFPNWSSTG	network server storage space	IPCS K: drive—firewall logs, queued mail, and cache

15.10.1 Saving the Firewall

All objects related to the firewall must be saved. If you miss any object, the save is not usable for recovery. If the firewall is destroyed and the save is not usable, you have to install and configure the firewall manually. For this reason, keep the planning sheets used to create the firewall in a safe place. As changes are made, go back and record the changes on the planning sheet, or make notes and store them with the planning sheets.

To save the firewall, perform these tasks:

1. Save the AS/400 TCP/IP configuration information related to the firewall.
2. Create a library for the firewall backup save files.
3. Stop the firewall application.
4. Vary off the network server description (NWSD).
5. Save the firewall communications configuration objects.
6. Save the firewall configuration.
7. Save the firewall operational data.
8. Vary on the firewall NWSD.
9. Start the firewall application.

These steps are discussed further in the following section.

15.10.1.1 Saving the TCP/IP Configuration

Entries are added to the AS/400 TCP/IP configuration as part of the firewall install and configuration process. This information is stored in a set of files used by the AS/400 TCP/IP. Save these files as part of the normal save and restore process, and after any changes are made to the TCP/IP configuration of the AS/400 system. Refer to the *TCP Configuration and Reference*, SC41-5420, for details on saving all these files.

As an additional backup of the parts needed to restore the firewall, it may be easier to add this information back to the AS/400 TCP/IP configuration after you do the restore.

To record this information, complete these steps:

1. Type CFGTCP on the command line to view the Configure TCP/IP display.
2. Select Option 1 (Work with TCP/IP Interfaces) to view a list of all the TCP/IP interfaces defined to the AS/400 system.
3. Record the Internet address, subnet mask, and line description names associated with your firewall NWSD.

15.10.1.2 Creating a Library for the Firewall Back-up Files

In some of our examples, we use save files to store our firewall backups. You need to create a library and save file objects. As an alternative option, use tape media to store your backup.

15.10.1.3 Stopping the Firewall Application

To stop the firewall application, type the command:

```
ENDNWSAPP NWSAPP(*FIREWALL) NWS(firewall)
```

15.10.1.4 Varying Off the Firewall NWSD

To vary off the firewall NWSD, type the command:

```
VRYCFG CFGOBJ(firewall) CFGTYPE(*NWS) STATUS(*OFF)
```

15.10.1.5 Saving Firewall Communication Configuration Objects

First, save the communication objects used by the firewall. Table 39 on page 298 lists the lines, controllers, devices, and network server description used by the firewall. These objects provide the communication environment used by the firewall.

To save the firewall communication configuration objects, enter:

```
SAVCFG DEV(*SAVF) SAVF(firesave/commobj)
```

or

```
SAVCFG DEV(tape-device-name)
```

In the previous example, firesave is the save file library and commobj is the save file.

Note: After saving the firewall communication configuration objects, you must save the firewall configuration.

15.10.1.6 Saving Firewall Configuration

The firewall configuration, which includes filter rules, proxy policy, and host names, is in a server storage space. The server storage space is designated as the E: drive on the IPCS. The server storage space is located in the QUSRSYS library with the name firewall13, where firewall is the name of the NWSD.

To save the firewall configuration, enter either of these commands:

```
SAVOBJ OBJ(FIREWALL3) LIB(QUSRSYS) DEV(*SAVF)
        OBJTYPE(*SVRSTG) SAVF(FIREWALL/CONFIG)
```

or

```
SAVOBJ OBJ(FIREWALL3) LIB(QUSRSYS) DEV(tape-device-name)
        OBJTYPE(*SVRSTG)
```

Note: After saving the firewall configuration to a save file, save the firewall operational data.

15.10.1.7 Saving Firewall Operational Data

Firewall operational data, such as logs and queued mail, is kept in a network server storage space. The server storage space is designated as the K: drive. The network server storage space is in the /QFPNWSSTG directory with a name of firewall100 where firewall is the name of the firewall NWSD. The operational data contained on the K: drive is very dynamic. The reason you save this data is so you have a K: drive to attach to the firewall in the event of a restore.

To save the firewall operational data, enter either of these commands:

```
SAV DEV('/QSYS.LIB/FIRESAVE.LIB/OPER.FILE')
    OBJ('/QFPNWSSTG/FIREWALLOO)
```

or

```
SAV DEV('/QSYS.LIB/tape-device-name.DEVD')
    OBJ('/QFPNWSSTG/FIREWALLOO)
```

If you save to a save file and choose to save all the saved firewall data to tape, use the SAVSAVFDTA command. The SAVLIB command also saves these save files. For more information, refer to the *Backup and Recovery*, SC41-5304.

Note: After your save is complete, restart the firewall.

15.10.1.8 Varying on the Firewall Network Server Description

After you complete the save, you must vary on the network server description (NWSD). To vary on the firewall before you start it, enter:

```
VRYCFG CFGOBJ(firewall) CFGTYPE(*NWS) STATUS(*ON)
        RESET(*YES)
```

Note: Once you get the message *Network server FIREWALL is active*, start the firewall application.

15.10.2 Restoring the Firewall

If you experience a system loss or damage to the firewall objects, you need to restore the firewall. Table 39 on page 298 lists the objects involved and provides their location. Some of the objects used by the firewall are rebuilt each time you vary on the network server description. However, the firewall configuration, such as filter rules are not. These items must be restored from a user provided backup.

The restore of the firewall information must be done in a particular order because the server spaces must link to the network server description when the NWSD is restored to the system. If the server spaces are missing or damaged, the links cannot complete correctly. We recommend following the order presented in this section to restore your firewall environment.

To restore the firewall, perform these tasks:

1. Restore the firewall operational data.
2. Restore the firewall communications configuration objects.
3. Restore the firewall configuration data.
4. Add the AS/400 TCP/IP definitions required for the firewall communication.
5. Vary on the network server description.
6. Start the firewall application.

These steps are further discussed below.

15.10.2.1 Restoring Firewall Operational Data

Firewall operational data, such as logs and queued mail, is kept in a network server storage space. This server storage space is designated as the K: drive. The network server storage space is in the /QFPNWSSTG directory with a name of firewall100. This operational data is dynamic data. The reason you restore this data is so that you have a K: drive to attach to the firewall. After the restore, the K: drive contains mail and logs that were on the firewall at the time of the save. You must restore the firewall operational data before you restore the communications configuration objects used by the firewall. This allows the links to complete when the NWSD is restored.

To restore the firewall operational data, enter:

```
RST DEV('/QSYS.LIB/FIREWALL.LIB/OPER.FILE')
      OBJ('/QFPNWSSTG/FIREWALLOO')
```

or

```
RST DEV('/QSYS.LIB/tape-device-name.DEVD)
      OBJ('/QFPNWSSTG/FIREWALLOO')
```

Note: We assume that you successfully saved the operational data as discussed previously in Section 15.10.1.7, "Saving Firewall Operational Data" on page 300. After you successfully restore the firewall operation data, restore the firewall communications configuration objects.

15.10.2.2 Restoring the Firewall Communication Configuration Objects

Next, restore the communications objects used by the firewall. Table 39 on page 298 lists the lines, controllers, devices, and network server description used by the firewall. These objects provide the communications environment used by the firewall. When the network server description is restored as part of this process, the links to the storage spaces are also rebuilt. A new image of the IPCS C: drive is also created.

To restore the firewall communication configuration object, enter:

```
RSTCFG OBJ(FIREW*) DEV(*SAVF) OBJTYPE(*LIND *CTLD *DEVD *NWSD)
      SAVF(FIRESAVE/COMMOBJ)
```

or

```
RSTOBJ OBJ(FIREW*) DEV(tape-device-name)
OBJTYPE(*LIND *CTLD* DEVD *NWS)
```

Recommendation

Be selective in your restore. Specify that only firewall related objects are to be restored. Do *not* specify *ALL for the *OBJ* parameter. This prevents you from overwriting configuration information unrelated to the firewall that may have changed even though your firewall configuration has not.

Note: We assume that you successfully saved the firewall communication configuration objects as described in Section 15.10.1.5, “Saving Firewall Communication Configuration Objects” on page 299.

Your next step is to restore the firewall configuration objects.

15.10.2.3 Restoring the Firewall Configuration Data

Finally, we restore the firewall configuration data, which includes filter rules, proxy policy, and host names. All these objects are kept in a server storage space. This server storage space is designated as the E: drive on the IPCS. The server storage space is located in the QUSRSYS library with the name firewall3, where firewall is the name of the firewall NWS.

You must restore firewall configuration data before you restore the communication configuration objects used by the firewall. This allows the links to complete when the NWS is restored.

To restore the firewall configuration data, enter:

```
RSTOBJ OBJ(FIREWALL3) SAVLIB(QUSRSYS) DEV(*SAVF)
OBJTYPE(*SVRSTG) SAVF(FIRESAVE/CONFIG)
```

or

```
RSTOBJ OBJ(FIREWALL3) SAVLIB(QUSRSYS) DEV(tape-device-name)
OBJTYPE(*SVRSTG)
```

Note: We assume that you previously saved the firewall configuration data as described in Section 15.10.1.6, “Saving Firewall Configuration” on page 300.

For more information about the AS/400 system and firewalls, see *AS/400 Internet Security IBM Firewall for AS/400*, SG24-2162.

15.10.3 Saving and Restoring the Filter Rules Using the COPY Command

There are other specific objects, such as filter rules, stored on the IPCS. These are well documented in the AS/400 Firewall redbook referenced in Section 15.10.2.3, “Restoring the Firewall Configuration Data.” We do not discuss the save and restore of these objects in this redbook.

Chapter 16. Database Protection and Availability

Of all information center resources, data is probably the most important. Other resources, such as hardware, vendor software, and building facilities are all ultimately replaceable—most data is not. Data is also the most volatile and complex of all information systems sources and the most critical to the business.

The complexity and volatility of data makes it the most difficult resource to manage during recovery and requires an ongoing management process. Therefore, one of the key components of a recovery plan is a database protection strategy. That is, how to protect the data on the system, what data to back up, how often, and how to recover it. Components of availability affecting the AS/400 database are covered in this chapter.

Databases contain the crux of business information. Without information, managers cannot make decisions, users cannot perform their work, all employees work at a disadvantage. You can protect this information by ensuring that databases are:

- Backed up for off-site storage
- Secured against unauthorized access
- Designed for efficient and effective recovery

Note: System Managed Access Path Protection (SMAPP) is a database protection technique covered in Section 4.4, “System Managed Access Path Protection” on page 38.

The most important activity to protect databases is the design and planning involved to put everything back when recovery becomes necessary. This chapter outlines the following planning and design considerations for database protection:

- Database journaling
- Journaling of access paths
- Saving access paths
- Recovering access paths
- System jobs affecting database availability
- DB2 Multi-system Protection

16.1 Saving Database Files for Recovery

When saving database files, a number of options are available to you. Your choice depends partly on:

1. The level of protection you want to achieve
2. The amount of time you want to spend on your backup compared to the time the related restore operation will take

With DB2/400 you can:

- Save and restore just one file
- Save and restore the whole database
- Perform a partial recovery
- Recover to a point-in-time
- Use referential integrity to link several files together in a network

With this flexibility comes additional considerations to ensure the database is properly saved. The remainder of this chapter addresses these issues.

16.2 Logical and Physical Files in Different Libraries

Make sure you understand your database networks in detail. By the term “database networks,” we refer to a network of physical and logical files residing in various libraries on the same system. If logical files and physical files are in different libraries, make sure that the restore of physical and logical files is done in the right sequence. It is not enough to specify *NONSYS or *ALLUSR on the restore commands. The related files may not be restored in the correct order.

Libraries are restored in alphabetic order. This causes problems if the logical files are stored in libraries named with letters “early in the alphabet” to their corresponding physical files in a “later in the alphabet” named library. As every environment is different, there is no standard method of restoring. We describe tools that can help you decide how to do your restores.

Note: Access paths are saved when you save the physical files because the physical file contains the data that is associated with the access path. Access paths are not saved when the physical and logical files are in different libraries. No message is issued to warn you that access paths are not saved in this situation. When you save the logical file, you save only the description of the logical file, which is all that is needed for a proper recovery.

An understanding of your database network can be obtained with the commands:

- ANZDBF
- WRKASP

The ANZDBF command is described in Section 16.3, “ANZDBF Command.” The WRKASP command is described in Chapter 9, “Tools for Automating System Management Functions” on page 109.

16.3 ANZDBF Command

The Analyze Database File (ANZDBF) command is available with the Performance Tools product. It is used to produce two reports – a physical-to-logical file relationship and a logical-to-physical file relationship report for files in a given library. The information is also saved in member QAPTAZDR of the database file QPFRDATA/QAPTAZDR. It can then be used as an input to the Analyze Database File Keys (ANZDBFKEY) command or user programs for further analysis.

The following steps explain how to use the Analyze Database File (ANZDBF) command against a library. The library used in our example is EDALIB:

1. Type ANZDBF.
2. Press F4.
3. Enter EDALIB for the name of the library that you want analyzed, as shown in the following figure.

```

Analyze Database Files (ANZDBF)

Type choices, press Enter.

Application libraries . . . . . > EDALIB      Name
      + for more values
Job name . . . . . ANZDBF      Name, ANZDBF
Job description . . . . . QPFRJOB      Name, *NONE
Library . . . . . *LIBL      Name, *LIBL, *CURLIB

```

Figure 93. Analyze Database Files (ANZDBF)

After entering EDALIB, a list of files appears, each with their own list of dependent files, as shown in Figure 94.

```

12/08/97 13:49:02      Database Relation Cross Reference
Type
P=Phy
L=Lgl  File           Library      Depnd  Dependent  Dependent  Depncy
      File           Library      Count  File       Library    Type
      -----          -
P      ADHOC00015      SYSCODW      1      ADHOC00016  SYSCODW      D
      CONV_00001
      CUST_00001      SYSCODW      2      CUST_00002  SYSCODW      D
      CUST_00003      SYSCODW      2      CUST_00003  SYSCODW      D
      CUST_00004      SYSCODW      2      CUST_00005  SYSCODW      D
      CUST_00006      SYSCODW      2      CUST_00006  SYSCODW      D
      CUST000001      SYSCODW      1      CUST000002  SYSCODW      D
      EIS_E00001      SYSCODW      1      EIS_E00002  SYSCODW      D
      EIS_F00001      SYSCODW      2      EIS_F00002  SYSCODW      D
      EIS_F00003      SYSCODW      2      EIS_F00003  SYSCODW      D
      EIS_STATUS

```

Figure 94. Database Relation Cross Reference

A "D" in the *Depncy Type D/A* column indicates a file with data dependencies. An "A" indicates a file with an access path shared with a second access path.

The data collected by running ANZDBF can be further analyzed with the ANZDBFKEY command.

The ANZDBFKEY command produces an Analysis of Keys for Database report, as shown in Figure 95 and a *Key Fields and Select/Omit Listing* report as shown in Figure 96 on page 306. These reports help you understand how your database network is organized.

```

12/08/97 14:02:58      Analysis of Keys for Database
Physical File OPNBB1PF  Library SYSCODW
File      Library  Logical  Maint      * Key Fields  Major to Minor * No.
      Format      Key Fields
OPNBB00004 SYSCODW  OPNBB1PF  I  WKNO  AYR      CUNO      3
OPNBB00005 SYSCODW  OPNBB1PF  I  ORIN  LN03     ITEM      3
OPNBB00001 SYSCODW  OPNBB1PF  I  ORIN  LN03      2
OPNBB00002 SYSCODW  OPNBB1PF  I  ITEM  CPVN      2
OPNBB00003 SYSCODW  OPNBB1PF  I  TSDT  CUNO      2

```

Figure 95. Analysis of Files per Library

```

12/08/97 14:02:58      Key Fields and Select/Omit Listing

  File      Library      Order  Path Type Unique Maintenance
PHY OPNBB1PF SYSCODW      FIFO  ARRIVAL

  Based on      Format      Key Field Seq Sign Zone Alt
  File      Library      Order  Path Type Unique Maintenance
LGL OPNBB00004 SYSCODW      FIFO  KEYED      N  *IMMED
  Based on      Format      Key Field Seq Sign Zone Alt
  OPNBB1PF SYSCODW OPNBB1PF  WKNO      SIGN
  AYR      SIGN
  CUNO

  File      Library      Order  Path Type Unique Maintenance
LGL OPNBB00005 SYSCODW      FIFO  KEYED      Y  *IMMED
  Based on      Format      Key Field Seq Sign Zone Alt
  OPNBB1PF SYSCODW OPNBB1PF  ORIN
  LN03      SIGN
  ITEM

```

Figure 96. Key Fields and Select/Omit Listing

These reports show file dependencies and key field characteristics that are useful for system administrators to understand the database network.

16.4 Referential Integrity Save and Restore Considerations

If you have a database structure with parent and dependent files, make sure these files are saved together. They should also be restored together. Otherwise, DB2/400 will detect inconsistencies between the parent and dependent files and put them into a check pending status. The check pending status must be resolved before users and applications can use the files.

For a thorough discussion of referential integrity and save and restore considerations, please refer to the redbook *Database2/400 Advanced Database Functions*, SG24-4249.

16.5 Save and Restore Tips for Trigger Programs

The Add Physical File Trigger (ADDPFTRG) command adds a trigger to call a trigger program to a specified physical file. A trigger is:

1. An event definition stored in a physical file that calls a trigger program when a specified operation is issued on the physical file
2. An exit program, called by a trigger, that contains a set of trigger statements

Save the trigger program with the database objects. If this is not done and the files are restored, errors may occur when you try to use the database again. This goes against the advice of keeping data and programs apart, but is a recommended if you use triggers.

A change operation can be an insert, update, or delete. The trigger program can be called before or after a change operation occurs. Thus, we have a possibility to affect six different trigger programs for each physical file.

An exclusive-no-read lock is held on the physical file when adding a trigger to that file. All logical files built over the physical file are also held with an exclusive-no-read lock.

Once a trigger is added to the physical file, all members of that specified file are affected by the trigger. When a change operation occurs on a member of the specified file (as an update, insert, or delete operation), the trigger program is called. The trigger program is also called when a change operation occurs on either a dependent logical file or a Structured Query Language (SQL) view that is built over the physical file.

Save and restore functions do not search a database file for a trigger program during the time of the save nor restore. It is the user's responsibility to manage the program.

During run time, if the trigger program has not been restored, an error results that halts processing of the program. The trigger program name, the physical file name, and the trigger event information is returned.

If the trigger program is restored to a different library, the change operation fails because the program is not found in the original library. An error results that halts processing of the restore. The trigger program name, the physical file name, and the trigger event information is returned in the error message.

There are two ways to recover in this situation:

- Restore the trigger program to the library in which it originally resided.
- Create a new trigger program with the same name in the new library.

Find more information on trigger programs in *DB2 for AS/400 Database Programming*, SC41-5701, and *Database2/400 Advanced Database Functions*, SG24-4249.

16.6 Save and Restore Relational Database Directories

The relational database (RDB) directory is made up of files that are opened by the system at IPL time. As such, the RDB is not an AS/400 object. Consequently, you cannot use the SAVOBJ command to directly save these files. Save the RDB by creating an output file from the relational database directory data. You can use this output file to add entries to the directory if the directory becomes damaged (similar to how the Retrieve Configuration Source (RTVCFGSRV) command is used to recreate configuration objects).

When entries have been added to the RDB and you want to save the RDB, specify the OUTFILE parameter on the Display Database Directory Entry (DSPRDBDIRE) command to send the results to an output file. The output file can be saved to tape, diskette, or a save file and restored to the system.

To illustrate, the following example restores RDB entries for the Survey Limited Corporation OUTFILE named SLCRDB created by:

```
DSPRDBDIRE RDB(*ALL) OUTPUT(*OUTFILE) OUTFILE(SLCRDB)
```

The sample CL program that follows applies to V4R2 systems. The CL program reads the contents of the SLCRDB output file and adds RDB directory entries using the Add Relational Database Directory Entry (ADDRDBDIRE) command.

```

/*****
/* - Restore RDB Entries from output file created with:      - */
/* - DSPRDBDIRE OUTPUT(*OUTFILE) OUTFILE(SLCRDB)            - */
/* -      from a V4R2 or later level of OS/400              - */
/*****
PGM
DCLF FILE(SLCRDB)                      /* See prolog concerning this */

/* Declare Entry Types Variables to Compare with &RWTYPE    */
DCL &LOCAL   *CHAR 1
DCL &SNA     *CHAR 1
DCL &IP      *CHAR 1
DCL &ARD     *CHAR 1
DCL &ARDSNA  *CHAR 1
DCL &ARDIP   *CHAR 1

/* Initialize Entry Type Variables to Assigned Values        */
CHGVAR &LOCAL '0' /* Local RDB (one per system) */
CHGVAR &SNA   '1' /* APPC entry (no ARD pgm) */
CHGVAR &IP    '2' /* TCP/IP entry (no ARD pgm) */
CHGVAR &ARD   '3' /* ARD pgm w/o comm parms */
CHGVAR &ARDSNA '4' /* ARD pgm with APPC parms */
CHGVAR &ARDIP '5' /* ARD pgm with TCP/IP parms */

RMVRDBDIRE   RDB(*ALL) /* Clear out directory */

NEXTENT: /* Start of processing loop */
RCVF /* Get a directory entry */
MONMSG MSGID(CPF0864) EXEC(DO) /* End of file processing */
  QSYS/RCVMSG PGMQ(*SAME (*)) MSGTYPE(*EXCP) RMV(*YES) MSGQ(*PGMQ)
  GOTO CMDLBL(LASTENT)
ENDDO
/* Process entry based on type code */
IF (&RWTYPE = &LOCAL) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) )

ELSE IF (&RWTYPE = &SNA) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +
  DEV(&RWDEV) LCLLOCNAME(&RWLLOC) +
  RMTNETID(&RWNTID) MODE(&RWMODE) TNSPGM(&RWTPN) )

ELSE IF (&RWTYPE = &IP) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWSLOC *IP) +
  TEXT(&RWTEXT) PORT(&RWPORT) )

ELSE IF (&RWTYPE = &ARD) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +
  ARDPGM(&RWDLIB/&RWDPGM) )
ELSE IF (&RWTYPE = &ARDSNA) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +
  DEV(&RWDEV) LCLLOCNAME(&RWLLOC) +
  RMTNETID(&RWNTID) MODE(&RWMODE) TNSPGM(&RWTPN) +
  ARDPGM(&RWDLIB/&RWDPGM) )

ELSE IF (&RWTYPE = &ARDIP) THEN( +
  QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWSLOC *IP) +
  TEXT(&RWTEXT) PORT(&RWPORT) +
  ARDPGM(&RWDLIB/&RWDPGM) )

```

```
GOTO CMDLBL(NEXTENT)
```

```
LASTENT:  
RETURN  
ENDPGM
```

The following example shows the same program for systems running V3R1 through V4R1.

```
/* *** Restore RDB Entries from output file created with:  
PGM  
DCLF FILE(SLCRDB)  
RMVRDBDIRE RDB(*ALL)  
NEXTENT:  
RCVF  
MONMSG MSGID(CPF0864) EXEC(DO)  
QSYS/RCVMSG PGMQ(*SAME (*)) MSGTYPE(*EXCP) RMV(*YES) MSGQ(*PGMQ)  
GOTO CMDLBL(LASTENT)  
ENDDO  
IF (&RWRLOC = '*LOCAL') DO  
QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT)  
ENDDO  
ELSE IF (&RWRLOC = '*ARDPGM') DO  
QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +  
ARDPGM(&RWDLIB/&RWDPGM)  
ENDDO  
ELSE IF (&RWDPGM *NE ' ') DO  
QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +  
DEV(&RWDEV) LCLLOCNAME(&RWLLOC) RMTNETID(&RWNTID) +  
MODE(&RWMODE) TNSPGM(&RWTPN) +  
ARDPGM(&RWDLIB/&RWDPGM)  
ENDDO  
ELSE DO  
QSYS/ADDRDBDIRE RDB(&RWRDB) RMTLOCNAME(&RWRLOC) TEXT(&RWTEXT) +  
DEV(&RWDEV) LCLLOCNAME(&RWLLOC) RMTNETID(&RWNTID)  
MODE(&RWMODE) TNSPGM(&RWTPN)  
ENDDO  
GOTO CMDLBL(NEXTENT)  
LASTENT:  
RETURN  
ENDPGM
```

The files that make up the relational database directory are saved when the SAVSYS command is run. The physical file containing the relational database directory can be restored from the save media to your library using the following Restore Object (RSTOBJ) command:

```
RSTOBJ OBJ(QADBXRDBD) SAVLIB(QSYS)  
DEV(TAP01) OBJTYPE(*FILE)  
LABEL(Qnnnnnnvrmxx0003)  
RSTLIB(your-lib)
```

In this example, the relational database directory is restored from tape. The characters nnnnnn in the LABEL parameter represent the product code of Operating System/400 (for example, 5769SS1 for Version 4 Release 2). The vrm in the LABEL parameter is the version, release, and modification level of OS/400. The xx in the LABEL parameter is the last two digits of the current system language value. For example, 2924 is for the English language; therefore, the value of xx is 24. After you restore this file to your library, you can use the information in the file to recreate the relational database directory.

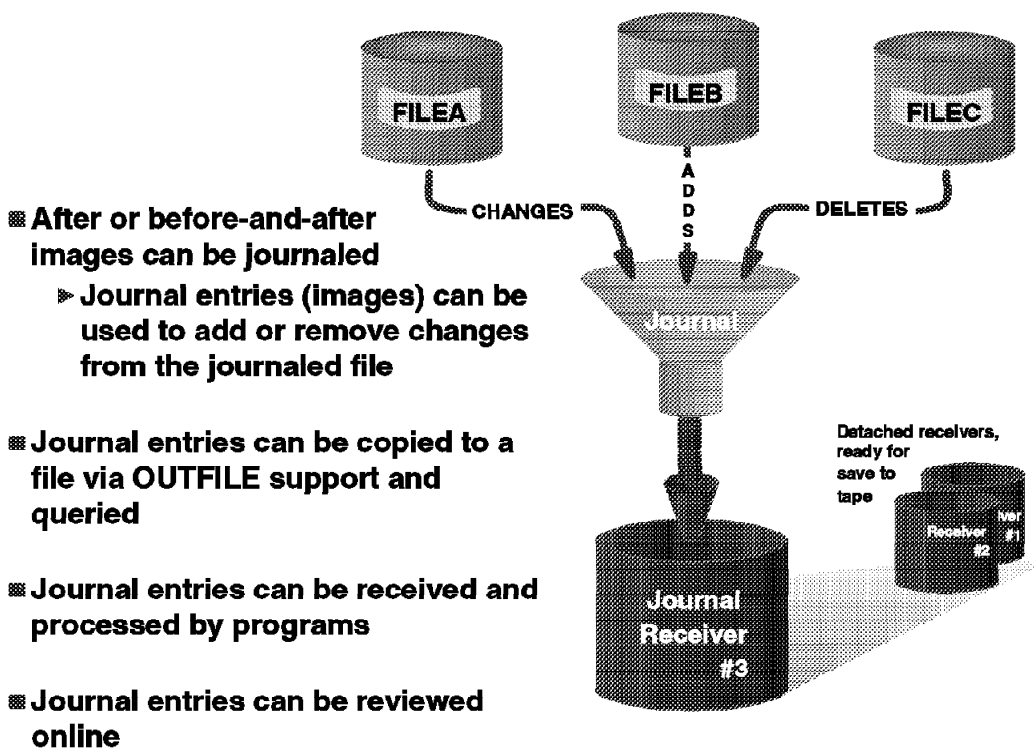
16.7 Stored Procedures

In V4R2, save and restore update system tables, views, sysprocs, and system parameters. If the stored procedure is not saved with the associated database objects, errors occur when the file is used again after the restore. Keep the program and database object in the same library to avoid this situation.

Refer to *Database2/400 Advanced Database Functions*, SG24-4249, for more information.

16.8 Database Journaling

Journals define which files and access paths to protect with journal management. This is called journaling a file or an access path. A *journal receiver* contains *journal entries* that the system adds when events occur that are journaled, such as changes to database files. This process is pictured in Figure 97.



SPOW002.PRZ-1

06/98

Figure 97. Database Journaling

Use journal management to recover changes to database files that have occurred since your last complete save.

Journaling is designed to prevent transactions from being lost if your system ends abnormally or has to be recovered. Journal management can also assist in recovery after user errors. Journaling allows the capability to roll back from an error to a stage prior to when the error occurred if both the before and after images are journaled. To recover, the restore must be done in the proper order:

1. Restore your backup from tape.
2. Apply journaled changes.

We recommend that you keep a record of which files are journaled. Refer to Section 16.9, “Determining Whether to Apply Journal Changes” for methods to track this.

A thorough discussion of journaling is beyond the scope of this chapter. For more information about journaling, refer to *Backup and Recovery*, SC41-5304.

16.9 Determining Whether to Apply Journal Changes

You may set up journaling yourself or use applications that use journaling. Two IBM applications which utilize journals are OfficeVision/400 and Client Access/400. Each of these use the QAOSDIAJRN journal in QUSRSYS. In addition, some applications provided by software vendors use journaling.

After a system failure and after databases are restored, some journal management may be required. You may have to apply the changes made to files since the last backup. The system manages journaling for OfficeVision/400 and Client Access/400 program with the QAOSDIAJRN journal. Check with your software vendor for recommendations on what management is required for journals they use.

Recommendation

Document the status of all journals and receivers prior to and following all system upgrades and software maintenance.

As part of any backup and recovery plan, maintain a list of journals and related journal receivers. To obtain a list of journals with related receivers, here are five options:

1. Enter the command:

```
DSPOBJD OBJ(*ALL/*ALL) OBJTYPE(*JRN) OUTPUT(*OUTFILE)
```

Then, create a CL program to process the output file created. Enter the command:

```
WRKJRNA JRN(journal-name) OUTPUT(*PRINT)
```

The entry `journal-name` is the `*OUTFILE` produced in the `DSPOBJD` command.

The resulting report displays information to determine whether any physical files are journaled and whether any journal entries exist that are more recent than your current restored copies of the files.

2. Use the `WRKJRNA` command to determine the relationship of journals and journal receivers.
3. Use the `PRTSYSINF` command as described in Section 9.5, “Print System Information Tool” on page 123.

- Use the following to produce a file that can then be used with Query/400 to determine which files are journaled and to which journal.

```
DSPFD FILE(*ALL/*ALL) TYPE(*MBR) OUTPUT(*OUTFILE)
      OUTFILE(QTEMP/FILELIST)
```

An advantage to this method is that the report shows files that used to be journaled but are no longer being journaled. The journal status is listed in addition to fields for the journal name and library. The output shows which images are being journaled.

- Use the following program to produce a database file containing a list of all journals on the system. The program prints a report for each journal, including a list of files associated with the journal and other useful information.

Note: To compile this program, you must first issue the DSPOBJD command to produce an outfile that the compiler can use to obtain field descriptions from.

```
      PGM
      DCLF      FILE(QTEMP/TEMPFILE)
      DSPOBJD   OBJ(*ALL/*ALL) OBJTYPE(*JRN) +
              OUTPUT(*OUTFILE) OUTFILE(QTEMP/TEMPFILE)
READ:   RCVF
        MONMSG  MSGID(CPF0864) EXEC(GOTO CMDLBL(ENDOFPGM))
        WRKJRNA JRN(&ODLBNM/&ODOBNM) OUTPUT(*PRINT)
        GOTO    CMDLBL(READ)
ENDOFPGM: ENDPGM
```

16.10 Considerations for SAVCHGOBJ when Journaling is Active

To use the APYJRNCHG command, the file must be journaled when it is saved. Save the file after journaling is started. **Do not** end journaling before saving the file with SAVCHGOBJ, otherwise information required for recovery is not saved.

Recommendation

Use the SAVCHGOBJ command when journaling is active.

Also be aware that the SAVCHGOBJ command defaults to **not** save journals. If you ever need to restore from the set of backup created when journaling was not active, the journals need to be recreated. Consider changing the OBJJRN parameter default to *YES for the SAVCHGOBJ command.

16.11 Database Journaling Performance

Database journaling has tradeoffs in system performance. Database journaling involves system overhead, but saves significant human overhead when recovery is needed. Some of these considerations are highlighted in this section.

Changes to a journaled file are written immediately to the journal receiver in auxiliary storage. This increases the disk activity on your system that can affect the overall system performance. Journaling also adds overhead to the resources involved in opening and closing files.

The larger the number of files being journaled and the greater the amount of database activity, the greater the impact on system performance. It can also

take longer to perform an IPL, especially if your system ends abnormally. If access paths are journaled, IPL time is reduced. Refer to Chapter 4, “IPL Improvements for Availability” on page 33, for information on other factors that affect IPL time.

16.11.1 Performance Tips for Journaling

To improve the performance for journaling activity, consider the following options:

- Isolate journal receivers in a user ASP. This reduces contention when accessing the disks.
- Place the fastest disks in the ASP used for journal receivers. In general, it is better to have multiple smaller capacity disks than one larger disk, because the number of disk arms is significant.
- Do not set the force-write ratio (FRCRATIO) parameter for physical files that are being journaled. Journal support provides better protection and performance than using a force ratio.
- Consider using record blocking when a program processes a journaled file sequentially. When you add or insert records to the file, the records are written to the journal receiver when the block is filled.
- Perform a commit operation approximately every 100 records to speed up batch jobs.
- The journal receiver size parameter with a remove internal entry value (RCVSIZOPT(*RMVINTENT)) on the Change Journal (CHGJRN) and Create Journal (CRTJRN) commands helps keep journal receivers size at a minimum. The RCVSIZOPT parameter indicates whether the system continuously removes internal entries as soon as they are no longer needed for recovery. The entries removed are primarily those needed for recovering access paths and internal control blocks. They are the hidden journal entries that consume sequence numbers.

By using the RCVSIZOPT parameter, the journal receiver size is decreased. Saving them to tape takes less space and time.

- Minimize the fixed-length portion of the journal entry by using the RCVSIZOPT(*MINFIXLEN) parameter on the CHGJRN and CRTJRN operations. The job, program, and user profile information are not deposited as part of the journal entry. This saves space on the journal receiver and reduces the overhead of making the journal deposit.
- Do not use DLTRCV(*YES) unless the journal is only used for access path protection or commitment control. Receivers are necessary to recover from a disaster.
- Let the system manage the change of journals.

Use the manage receiver parameter (MNGRCV(*SYSTEM)) on the Change Journal (CHGJRN) and Create Journal (CRTJRN) commands to attach a new journal receiver faster and with less system disruption. The MNGRCV parameter indicates whether the system is allowed to automatically initiate creating and attaching a new journal receiver when the attached receiver reaches its threshold value. This option provides easier journal receiver management.

In addition, MNGRCV(*SYSTEM) improves the I/O performance of a CHGJRN operation— for recovery purposes. Internally, the system records the

identity of the oldest journal entry to be read and processed on behalf of journaled objects that remain in main memory. During a CHGJRN operation for a user-managed journal, all journaled objects are forced from main storage to disk so that the synchronization point for an object is transferred to the new receiver. This can yield a large spike of disk-write activity on behalf of the CHGJRN operation.

By contrast, when the MNGRCV(*SYSTEM) parameter is specified for a journal, the spike of I/O activity no longer occurs because synchronization points are allowed to move naturally to the newly attached receiver. The system forces the journaled objects during normal system activity.

When MNGRCV(*SYSTEM) is specified for a journal, you can choose to:

- Let the system manage creation and attachment of new receivers.
- Manually initiate CHGJRN operations on fixed time intervals.

Both techniques take advantage of the reduced I/O. Therefore, you do not need to give up your current CHGJRN practices to gain the benefit. You can accomplish this by specifying less than the maximum threshold value for the journal receiver (for example, 1.5GB) together with MNGRCV(*SYSTEM) so that the system effectively only initiates a CHGJRN operation when a receiver reaches full status. The system managed attribute becomes your safety net, allowing you to employ your current strategy with faster performance. It kicks in only if your current journal management strategy gets “surprised” by a sudden burst of journal deposits.

When journals are system managed, the receiver is changed and sequence numbers are reset every IPL. With MNGRCV(*SYSTEM) specified, the system resets the sequence number of journal entries to “1” when it changes the journal.

Recommendation

If your application provider uses or depends on journaling, check with the provider prior to using system managed journals to make sure receiver sequence number resets are managed appropriately.

16.12 PTFs for CHGJRN performance improvements:

A number of PTFs are available to help improve the performance of the CHGJRN command, especially if you have a significant number of objects journaled to the same journal. These performance improvements are included in the base operating system beginning with V4R1. For previous releases, check out these PTFs:

- V3R1M0
MF18082 and SF39021
- V3R2M0
MF16873 and SF36298
- V3R6M0
MF17562 and SF46569
- V3R7M0

Note: These PTF numbers may have been superseded. Contact IBM support to make sure that you have the latest PTFs available.

16.13 Elimination of Lock Conflicts Between CHGJRN and RCVJRNE

Have you tried to do your Change Journal (CHGJRN) only to find that a lock conflict prevented the change from occurring because another job was doing a Receive Journal Entry (RCVJRNE)? Beginning in V3R2 and V3R6, improvements were made so that CHGJRN and RCVJRNE lock conflicts no longer exists. This makes it easier for you (or the system) to use CHGJRN during your normal business day.

16.14 Considerations for 1TB Maximum Access Path Size

On V4R1 systems and later, the default value for access path size for the Create Physical File and Create Logical File commands (CRTPF and CRTLF) is 1 terabyte (TB). On prior systems the default is 4GB. This means that instead of a three-byte index, a four-byte index is built with the implicit ability of granularity on the index.

The four-byte index circumvents a lot of seize conflicts. However, this change may have an impact on the implicit access path sharing for database files created prior to V4R1.

If you have an existing database with all three-byte indexes and you create or recreate a file with an index, that index becomes a four-byte index by default. It does not have the ability to implicitly share an existing access path of a three-byte index.

From a functional point of view, you cannot see any difference since the application runs in the same way. However, it can impact overall system performance when more implicit shares are removed and eventually a large number of access paths exist for both the three and four-byte versions.

Recommendation

Do not mix three- and four-byte access paths.

There are three approaches to avoid mixing three and four-byte access paths:

1. Leave the Access Path Size parameter (ACCPTHISIZ) default to *MAX1TB and recreate the database so all access paths are in four-byte mode. The downfall of this approach is that recreating the database can be time consuming.
2. Change the ACCPTHISIZ default back to *MAX4GB. This does not avoid potential seize conflicts, however.
3. Use DSPFD and DSPDBR commands for the critical files (those with large access paths and/or heavily used) to keep the files in synch with access paths. Migrate to four-byte indexes in a phased manner.

Use the CHGPF command to change the access path size of a physical file and the CHGLF command to change the access path size of a logical file.

Note: On systems prior to V4R2, the size of access paths for logical files is changed by deleting and recreating the logical file.

16.15 Journaling of Access Paths

An access path can be considered a definition of the order in which records in a database file are processed. A file can have multiple access paths.

Journal management can help keep a record of changes to access paths. This reduces the amount of time it takes the system to perform an IPL after an abnormal end. If your system abnormally terminates when access paths are in use, the system may have to rebuild the access paths before the files can be used again. This is a very time consuming process that can take many hours on a large, busy AS/400 system.

Two methods of access path protection are:

1. SMAPP (System Managed Access Path Protection)
2. Explicit journaling of access paths

Use SMAPP so the system can decide which access paths to protect, based on an overall target recovery time for access paths. Or, use explicit journaling for the specific access paths used for your most critical business applications.

SMAPP is further described in Chapter 4, "IPL Improvements for Availability" on page 33.

Note: You can use both SMAPP and explicit journals on the same system. For example, use SMAPP for interactive work during the day while files are being used. Then, use overnight journaling on the access path. Typically files used by batch runs are large and are not used for interactive jobs. Their access paths can be rebuilt manually during off hours when a failure occurs.

16.15.1 Access Path Journals Compared to SMAPP

There are advantages to using SMAPP over explicitly managing access path journals as depicted in the list below.

	Explicit Journals	SMAPP
Application knowledge required	Yes	No
Management of journal receivers required	Yes	No
Disk space consumption	Yes	No
Journal receiver placement	Yes	No
Enabling protection	Yes	No
Performance impact	Yes	No
Data integrity	Yes	No

- Application knowledge

Access-path journaling requires a thorough understanding of the application databases and its dependencies to ensure that all current and new important access paths are journaled.

SMAPP is automatic and constantly studies all the access paths to ensure that it can provide protection to meet the target recovery time specified for the configured ASP.

- Journal receiver

Access-path journaling requires you to manage the receivers and regularly detach them for offline storage.

SMAPP creates and manages a hidden logging area within the system and does not require any user management.

- Disk space consumption

Access-path journaling requires you to detach journal receivers using the Change Journal (CHGJRN) command. Journal receivers can overflow an ASP and consume hundreds of megabytes per day.

The SMAPP logging (tracking changes to access paths) area is circular and consumes less space. It has 78 bytes less information (time, job, and program entries) to record compared to a journal entry.

- Journal-receiver placement

Access-path placement requires journal receiver placement on a user ASP to achieve the best possible performance

SMAPP spreads the logging area among the number of disk arms available, with affinity for arms with write cache. SMAPP bundles the writes until the package is 128K full (on RISC systems) and writes them in parallel to all disk arms

- Enabling protection

Access-path journaling requires customers to issue the Start Journal Access Path (STRJRNAP) and Start Journal Physical File (STRJRNPF) commands along with the necessary setup commands required to start journaling.

SMAPP enables automatically when V3R1 is loaded on the system. The default target recovery time that SMAPP aims for is 150 minutes. This time can be changed for each configured ASP using the Edit Recovery for Access Path (EDTRCYAP) command. The changes take place immediately.

Underlying physical files do not need need to be explicitly journaled, as journaling is handled by SMAPP. You can view the current settings with the Display Recovery for Access Paths (DSPRCYAP) command.

- Performance impact

Access-path journaling performs a synchronous write to the disk for each database operation (add, delete, and update).

SMAPP provides just as good protection with bundled asynchronous writes therefore, reducing the total number of writes and performance overhead compared to synchronous writes. Performance also depends on how the target recovery time is set. The shorter the recovery time, the greater the performance overhead due to SMAPP.

- Data integrity

Access-path journaling data is written to the journal receivers and during recovery, data can be recovered to a known point. However, recovery may take several hours, or in some cases, days if the access paths are not protected.

An IPL or restricted state condition is required to activate SMAPP if it is turned off. Specifying *NONE for the system access path recovery time prompt on the EDTRCYAP command is better than specifying *OFF since it does not require a restricted state to change from *NONE to a specific time.

When *NONE is specified, no recovery is in place, but estimates are made of the time to recover access paths during an IPL. When you see the potential time spent, you may decide that SMAPP is worth the small overhead it entails.

SMAPP's primary objective is to reduce the access path rebuild time during an abnormal IPL. There is a possibility that data entered seconds before an abnormal termination may be lost and require re-input. However, the time spent recovering the system after an abnormal termination is significantly reduced. This, of course, is controlled by the target recovery time set for SMAPP.

16.16 Saving Access Paths

Restoring a database is typically an emergency action where the operators are pressured to make the recovery process take as little time as possible as to minimize the impact on system availability. This section describes differences in recovery when SMAPP is or is not activated.

The following chart illustrates what processes are performed during a recovery of the system when SMAPP is not in effect. Note that rebuilding access paths takes the most time to recover a system when the access path is not protected.




Figure 98. How Long it Takes to Rebuild Access Paths

Recommendation

Save access paths, even if a save with access paths takes longer than saving your database without access paths.

Following this recommendation offers users quicker access to the system after a system restore when the QDBSRVnn jobs consume CPU performing the rebuild.

If you do not save access paths, they cannot be restored, and therefore, will be rebuilt at run time as a side-effect of the restore itself. Those access paths not protected with SMAPP or journaling are the recover parameter in the CRTPF, which is subject to a rebuild during the next IPL following a crash just when the pressure is highest to get the system recovered. Whether these non-protected access paths are rebuilt at IPL or when first used depends on the recover parameter on the Create Physical File (CRTPF) and Create Logical File (CRTLFL) commands. The options are:

- *NO*—The access path of the file is rebuilt when the file is opened
- *AFTIPL*—The access path of the file is rebuilt after IPL is completed
- *IPL*—The access path of the file is rebuilt during the IPL operation

To save access paths, make sure that:

- ACCPTH(*YES) is specified on the save command.
- All the physical files that the access paths are based upon are in the same library.
- The logical files are MAINT(*IMMED) or (*DLY).

Note: Be aware that a save of access paths is **not** the default for the save commands. A save of access paths *is* the default if you use the SAVE menu options.

Refer to Section 5.8, “Unattended Saves Using the SAVE Menu” on page 61 for a further description of the changes to the SAVE menu.

The impact of rebuilding at file open is depicted in Figure 99 on page 320, which identifies rebuild time assuming that SMAPP is not enabled. Imagine how users productivity is affected when they spend hours waiting before they can begin work after the system is recovered, while the access paths are rebuilt. Use the Edit Rebuild of Access Paths (EDTRBDAP) command to control the number of access paths to be rebuilt.

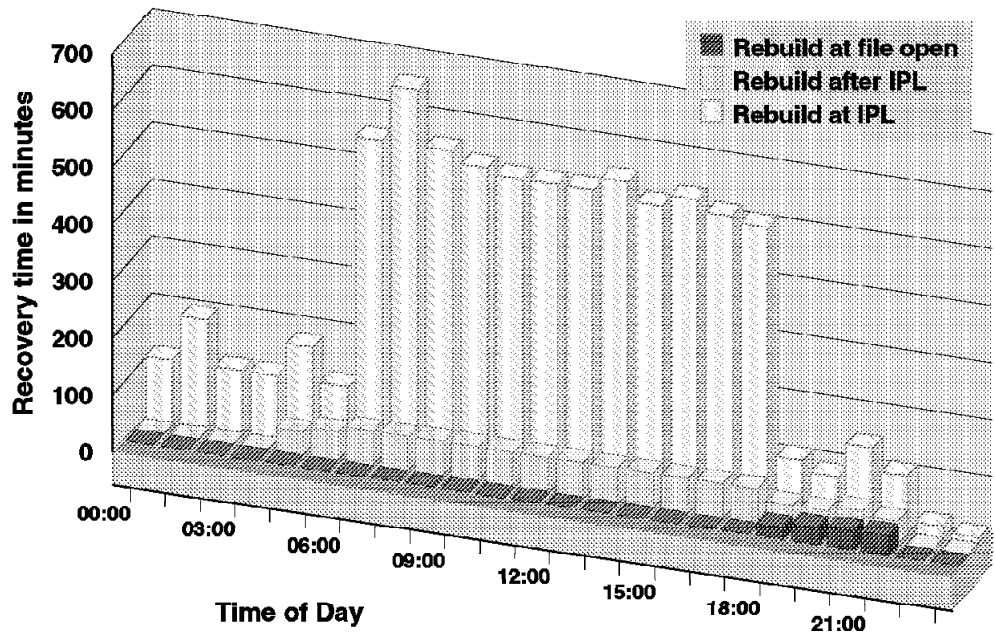


Figure 99. Access Path Recovery Exposures Based on Time of Day

Use the Change Command Default (CHGCMDDFT) command to change the command defaults for the SAVxxx commands to specify ACCPTH(*YES). Remember to use CHGCMDDFT after each release update, because the SAVxxx commands are replaced during the upgrade process and your defaults are overwritten.

Recommendation

Use the SAVE menu options to perform your saves. Access paths are saved when the SAVE menu is used to perform a backup.

Note: There are instances when access paths are rebuilt during the restore even though they are saved. For example, when an access path size of 1TB is restored to a system that does not support this size, the access paths are rebuilt to a size supported at that release.

16.17 Journal Entries Considerations for V4R2

Several changes affecting journal entries are described in this section. When using the Display Journal (DSPJRN), Retrieve Journal Entry (RTVJRNE), or the Receive Journal Entry (RCVJRNE) commands to view, retrieve, or receive journal entries, be aware of the following differences between V4R2 and prior release systems:

- Journalized object names

The name that appears for a journalized object is the same as it was at the time of the journal entry deposit. On releases prior to V4R2, the name that

appears for a journaled object is the name that the object was last known by on the system. For example, if you journal a file called OLDFILE, do a put of a record. Rename the file to NEWFILE. When you display the put record entry, NEWFILE is the object name—not OLDFILE. On V4R2, the name OLDFILE is displayed when viewing the entry.

- System name

The system name that appears in a journal entry is the system name where the journal entry was deposited. On previous releases, the system name that appears is the name of the system on which you performed the display, retrieve, or receive. On V4R2, if you restore a journal receiver from system A to system B, the system name appears for the journal entries in that receiver on system A, not system B as on previous release systems.

- Journal entries are no longer sent

The following two journal entries are not sent to the journal and cannot be viewed when looking at the journal entry data:

```
Journal Code - F   - for database file member operation
                Entry Type - PM   - access path moved
                Entry Type - PN   - access path renamed
```

16.18 Journal Receiver Protection

To protect your journal receivers, you can:

1. Mirror the ASP where the journal receivers reside.
2. Use RAID DASD for the ASP where the journal receivers reside.
3. Use remote journal support provided in V4R2.

For more information on mirroring and RAID protection refer to Section 3.2, “Device Parity Protection” on page 27. For more information on remote journal support, refer to Chapter 17, “Using Remote Journals to Improve Availability and Recovery” on page 327, and the *Backup and Recovery* manual, SC41-5304.

16.19 Multi-member Database File Save Performance

To enhance database save performance, V4R1 and later systems maintain a member list for files. Using a member list allows database file save functions to asynchronously bring database members into main storage, greatly reducing the time waiting for page faults. A form of parallelism occurs by bringing the members into main storage in parallel while actually starting the save process on the members.

Prior to V4R1, a file with multiple members is essentially a linked list. During a save process, each member is page faulted into main storage before storage management brings the next member into main storage. In other words, each member is synchronously brought into main storage before starting on the next member. This is one of the primary reasons why saving multiple member files is slower on systems prior to V4R1.

16.20 Database Server Jobs

QDBSRVnn and QDBSRVXR are system jobs for database functions. There is no direct way for users to control these server jobs. However, you can change the amount of work they are required to do. For example, you can:

1. Use the Edit Rebuild of Access Paths (EDTRBDAP) command to control the number of access paths to rebuild.
2. Increase the size of the *BASE pool to allow the jobs to get more work done since system functions use *BASE. You can also increase the size of the *BASE pool to allow the jobs to finish more work.

QDBSRV01 is the server job that handles events for database, commit, journal, and System Managed Access Path Protection (SMAPP) (for example). It may also perform some work for journal, database and commit operations. Its run priority is level nine.

QDBSRV02 and QDBSRV03 jobs run at a priority level of 16. These jobs also perform functions for database, journal, and commit operations. The normal job status is DEQW.

QDBSRV04 through QDBSRV05 jobs run at a priority level of 52. One or more of these jobs perform the same function that QDBSRV1 and QDBSRV2 did prior to Version 3 Release 1.

On four-way processors, QDBSRV02 through QDBSRV05 jobs run at a priority level of 16, and QDBSRV06 through QDBSRV09 jobs run at a priority level of 52.

Note: QDBSRV1 and QDBSRV2 are database background processes on systems prior to V3R1. They run at a priority level of 52 and rebuild the access path when logical files are restored to the system. These jobs run during an abnormal IPL or when a logical file's access path maintenance plan is changed from *DELAY or *REBUILD to *IMMED.

QDBSRVXR handles most of the functions for the system cross-reference files. QDBSRVXR runs at a priority level of 0 with a normal status of DEQW.

Cross-reference files are journaled in QRECOVERY/QDBJRN. Deletions, restores, and copies of a large number of files causes a large number of messages logged to the QSYSOPR message queue related to the changing of receivers.

16.21 DB2 Multisystem

The DB2 Multisystem for OS/400 is a separate component installed as Option 27 of OS/400. It has been available since V3R2 for AS/400 CISC systems and V3R7 for RISC systems.

The DB2 Multisystem feature provides a straight forward method to distribute your database across multiple AS/400 systems. Up to 32 AS/400 systems can be connected in an APPC network and become a single DB2 for OS/400 database. DB2 Multisystems support OptiConnect to provide the fastest communications link available over and above networks. Refer to Chapter 18, "OptiConnect for OS/400" on page 351, for more information about OptiConnect, including distance limitations.

A query can be initiated on any of the systems. The DB2 Multisystem broadcasts the query across the network and executes in parallel. Each of the participating systems execute the query on its own portion of the data.

Transparency is one of the major advantages provided by the DB2 Multisystem. Distributed files behave as local database files behave. Users run queries and applications accessing the files as if the data is local.

16.22 Restoring Distributed Files

As with save operations, restore operations only affect the system where they are performed. It is the (human) system administrator's (or manager's) responsibility to keep parts of the distributed files synchronized on different systems in a network.

If a system in the network is lost, you have two alternatives:

1. Replace the missing system with a new one. Configure the communications and relational database directory the same way. Then restore the portions of the distributed files onto this system and redistribute them if desired.
2. Restore all portions of each distributed file on a single system. For each distributed file, create a local copy containing all of the records. Once you rebuild the entire file, redistribute the records among the remaining systems if desired.

For a more detailed discussion of save and restore strategies for DB2 Multisystem files, refer to the redbook *Database Parallelism on the AS/400*, SG24-4826.

16.22.1 Distributed Files Backup Considerations

Spreading your database over multiple AS/400 systems raises some concerns relative to backup and recovery.

Each portion of a distributed file is an individual object. Therefore, journal management and save and restore operations on distributed files are performed on each individual system. In other words, to save or restore a distributed file, you need to process every portion of the file separately.

To save a distributed file, use the SAVLIB or SAVOBJ commands. You only save the portion of the entire distributed file located on the system where the save command is issued.

As a result of a save, you:

- End up with a set of tapes containing a consistent copy of all portions of a distributed file.
- Run the save from a single workplace rather than signing on to each system in the network.

The DB2 Multisystem feature provides a convenient solution for both of these requirements. Consider the following concepts:

- When you want to save an object, use the Allocate Object (ALCOBJ) command to place an exclusive lock and prevent anyone from modifying the object before the save operation is complete. The ALCOBJ command has a

global effect on a distributed file. You establish a lock on a distributed file by issuing the ALCOBJ once from any of the systems in the network.

To release the lock, use the Deallocate Object (DLCOBJ) command.

- A distributed file can be a target of the SBMRMTCMD command, as well as a DDM file.

When you issue the Submit Remote (SBMRMTCMD) command and refer to a distributed file in the DDMFILE parameter, the requested command is executed on every remote system for the distributed file.

If you journal your distributed files, you need to save all of the journal receivers on all participating systems simultaneously. This allow you to have a backup set reflecting a single consistent point of recovery across all systems.

16.23 Distributed File Back Up Scenario

This section outlines the steps to back up distributed files. In our example, we use the CUSTREC file in the CUSTLIB library.

To make a consistent save of a distributed file (in this case, file CUSTREC in library CUSTLIB) across all the participating systems, use the following procedure:

1. Place an exclusive lock on all parts of the distributed file. Enter:

```
ALCOBJ OBJ((library1-name/CUSTREC *FILE *EXCL))
```
2. Save the local part of the file to a backup tape. Enter:

```
SAVOBJ OBJ(CUSTREC) LIB(library1-name) DEV(tape-device-name)  
OBJTYPE(*FILE)
```
3. Save all of the remote parts of the file to respective backup tapes. It is assumed that operators have loaded the correct tapes on the various systems. For this, use the SBMRMTCMD command:

```
SBMRMTCMD  
  CMD(SAVOBJ OBJ(CUSTREC) LIB(library1-name) DEV(tape-device-name)  
  OBJTYPE(*FILE) DDMFILE(library1-name/CUSTREC)
```
4. Detach the current journal receiver and attach a new journal receiver on the local system:

```
CHGJRN JRN(library1-name/journal-name) JRNRCV(*GEN)
```
5. Perform the same procedure for the remote systems:

```
SBMRMTCMD  
  CMD(CHGJRN JRN(library1-name/journal-name) JRNRCV(*GEN)  
  DDMFILE (library1-name/CUSTREC)
```
6. Save the detached journal receiver on the local system:

```
SAVOBJ OBJ(journal-receiver-name) LIB(library2-name) DEV(tape-device-name)  
OBJTYPE(*JRNRCV)
```
7. Save the journal receivers you just detached on the remote systems:

```
SBMRMTCMD CMD(SAVOBJ OBJ(journal-receiver-name)  
  LIB(library2-name) DEV(tape-device-name) OBJTYPE(*JRNRCV))  
  DDMFILE(library1-name/CUSTREC)
```

8. Once saved, delete this journal receiver on the local system:

```
DLTJRNRCV JRNRCV(library2-name/journal-receiver-name)
```

9. Delete the journal receiver on the remote system:

```
SBMRMTCMD CMD(DLTJRNRCV  
JRNRCV(library2-name/journal-receiver-name))  
DDMFILE(library1-name/CUSTREC)
```

10. Release the lock from the file on all systems:

```
DLCOBJ OBJ(library1-name/CUSTREC *FILE *EXCL)
```

You now have a set of tapes (one for each system) that together hold a complete backup image of a distributed file.

Note that since update activity on different systems for the distributed file can vary, the number and names of journal receivers for all systems can be different. Track saved journal receivers separately for each system or develop a single consistent journal receiver management policy for all systems. In the previous example, journal receivers are changed simultaneously on all systems and are named the same. You can use a similar convention or devise your own.

16.24 Considerations for a Multinational Environment

Support is available on V4R2 systems to determine the decimal format for edited numeric output fields when a display or printer file is opened. This format is the default for files created on V4R2 systems using the Create Display File (CRTDSPF) or the Create Printer File (CRTPRTF) commands. Files created prior to V4R2 can be changed to use the national language decimal format by specifying DECFMT(*JOB) on the Change Display File (CHGDSPF) or the Change Printer File (CHGPRTF) command. Use CRTDSPF or CRTPRTF to recreate files that have three-character and four-character EDTCDE(Y) fields to get the national language decimal format support.

On systems prior to V4R2, the support for determining the decimal format when the file is opened is provided with PTFs. As part of the PTF support, you can create the QWSDECFMT data area to indicate the decimal format to use. The value in the data area overrides the value in the QDECFMT system value.

16.25 For More Information

For more information on database operations and management, refer to the URL:

<http://www.as400.ibm.com/>

Follow the *Software*, *Software Solutions—Database*, *Technical Information*, and *Tips* links for more information on database functions, including journaling tips.

Chapter 17. Using Remote Journals to Improve Availability and Recovery

The ability to have AS/400 database updates shared concurrently with a second AS/400 database enhances the availability strength of the AS/400 system. This database replication function can be implemented with application code that employs the remote journal function on V4R2 or later systems.

The remote journal function is part of the base OS/400 beginning with V4R2. Implementing remote journals enables the association of journals and journal receivers between source and target AS/400 systems. By creating an association of a local journal and its journal receivers to a remote journal and its associated receivers, the remote journal function enables journal receiver replication from a source system to a target system. When used in coordination with communications interfaces, such as OptiConnect for OS/400, SNA or TCP/IP, customers can create clustered environments that provide 24 x 365 availability and full system redundancy.

Remote journals help offload CPU consumption from the local system so it can achieve more throughput. Remote journals replace customer programming with a more efficient system programming interface to capture and transmit journal images between source and target systems.

High availability solutions from IBM, business partners and user written applications on systems prior to V4R2 employ the Receive Journal Entry (RCVJRNE) command on the source system to perform the replication process. An exit program receives journal entries on the source system and sends them to the target system using an available communication link. Remote journal support allows this RCVJRNE overhead to be shifted to the remote system.

Find additional information on high availability solutions from IBM and its business partners in Appendix E, "High Availability Solutions" on page 391, in this redbook. These solution providers handle many of the considerations discussed in this chapter.

17.1 Remote Journal Function

The remote journal function performs journal entry replication at the Licensed Internal Code (LIC) layer. Moving the replication to this lower layer:

- Shifts more of the replication overhead from the source system to the remote system, which improves overall system performance on the source (primary) system
- Improves performance for journal replication
- Allows the option to replicate synchronously with the operation causing journal database modification
- Moves the journal receiver save operations to the remote system, offloading the source (production) system

All methods on systems prior to V4R2 to get journal entries to a backup system require that you use a user-written apply program to apply these database changes on the replicated database on the backup system. Either a business

partner product or a written apply program is required to apply any changes to the replicated database on a backup (target) system.

Hot site backup and high availability applications are good examples of applications that can benefit from the use of remote journals. If a switch over to the target system is necessary, the business partner or user-written apply and remove program must remove partial commitment control transactions from the backup system. This must take place prior to using the backup (target) system as the primary (production) system. Likewise, once the source system is repaired, business partner software or a user written apply and remove program is required on V4R2 systems. Such programs replay the database changes back to the original database on the production system that were made on the target system while the production system was offline.

The remote journal function on V4R2 differs from methods used on systems prior to V4R2 for replicating data. The differences can be seen by comparing Figure 100 and Figure 101.

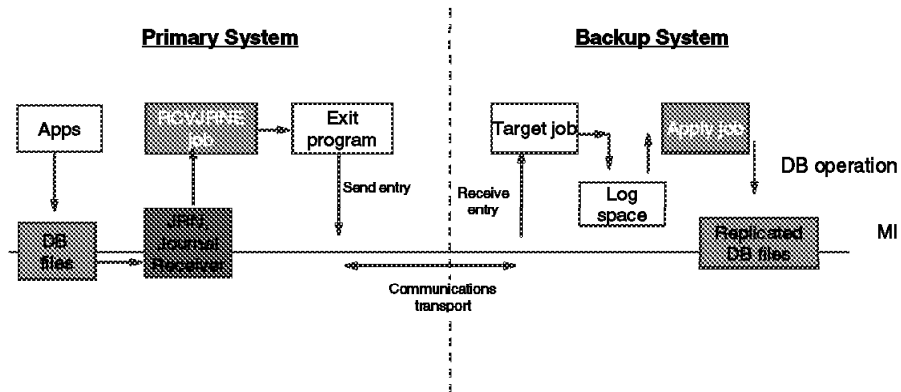


Figure 100. Replication without Remote Journal

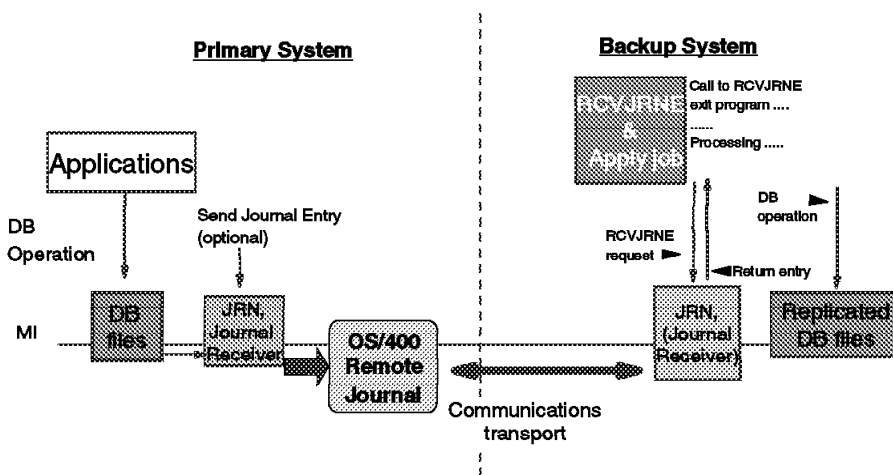


Figure 101. Replication with Remote Journal

Sending journal entries to a backup system does not mean that the journal entries are *only* being sent to a backup system. In all cases of using the remote journal function, sending journal entries to a journal that resides on the same system as the source database is always performed. Remote journals allow journal entries to be optionally sent to one or more additional backup systems. The journal entries sent to the backup system are deposited by the system into a duplicate journal receiver. From an external viewpoint, the duplicate journal receiver is identical to the receiver that is attached to the journal on the primary system. For recovery purposes, this allows receivers on the backup system to be used when needed for recovery of the primary system. Essentially, the receivers on the backup system can be viewed externally as “saved” versions of the receivers on the primary (production or source) system.

Note the following points:

- There is no application program required on the target system.
- The processing of journals on systems prior to V4R2 perform *above* the machine interface (MI) without remote journal support and *below* the MI with remote journal support on V4R2 systems.
- No exit program to time or manage the journal synchronization is required with a remote journal implementation.

The following differences enable more efficient journal functions:

- Lost and trapped transactions are prevented.
- There are less disk writes on the source system than with former journal operations.
- DASD is used more efficiently.
- There is less CPU overhead than prior support.
- CPU cycles are freed up on the production and source machine.
- Work is shifted to the backup or target machine.
- Database images can be sent to the target machine in real time.
- There is no delay in replication.

17.2 When the Remote Journal Function Can Be Used

The most common reasons for using the remote journal function include:

- For high availability (24 x 365)
- To free up CPU cycles on the source/production machine
- To prevent lost transactions
- When a quick and easy switch back to the recovered system is needed

Be sure that you are running the V4R2 refresh of any business partner product that incorporates the remote journal function (such as MIMIX, OMS Gold, or Transformation Server for AS/400). The application product refresh must occur on each machine that implements remote journals. That is, both source and target must step up to V4R2.

When remote journals are set up, take note of the following conditions:

- No application changes are required (unless you are building your own or using the DataPropagator product).
- There are no special features to install (unless you are using OptiConnect as the transport vehicle)

- No tuning is mandated (but may be of some benefit).
- The business partner or user needs to create and maintain relational databases directory entries on both the source and target systems.

There is extra maintenance of two receivers on two systems to save or delete receivers as appropriate.

17.3 Remote Journal Transport Protocol

With remote journals set up, journal entries can be transported from the source system to the target system in real time. All transport activities are managed below the MI level. The transport functions are performed from memory to memory and do not wait for updates to be written to disk on the target system. If the communications link between the source and the target systems fails, source applications continue to execute. When the remote journal function is re-established, all intervening updates are sent to the target system.

There are three supported transport mechanisms:

- OptiConnect for OS/400 (optical bus)
- TCP/IP
- SNA (APPC)

Relational database directories entries define the communications protocol and remote system.

Remote journals can broadcast to up to 255 secondaries or target journals. Cascading (sending from the secondaries to subsequent systems) can be done without limitation.

17.4 Remote Journal Replication Modes

Journal entries can be replicated to the target system from the source system either synchronously or asynchronously. You need to understand the advantages and disadvantages of either when you implement remote journals. They are discussed in the sections below.

17.4.1 Synchronous Delivery Mode

Synchronous delivery mode means that the application program that initiated the database change on the source system remains input inhibited. It remains in this state until the journal entry is present in main storage on the target system. The advantage of this approach is that the target system has all of the journal entries (database changes) as they are being made in real-time on the source system. Synchronous delivery allows for full recovery of the entire database on the target system when an outage is experienced on the source system.

Although sending journal entries synchronously to a target system has some impact to journal throughput on the source system, this impact is generally quite modest. The amount of impact depends on resource usage and the communications rate of making database changes.

17.4.2 Asynchronous Delivery Mode

Sending a journal asynchronously means the journal entry goes to the target system *after* control is returned to the end-user application that deposited the journal entry on the source system. From a recovery standpoint, the asynchronous delivery mode is less desirable because the target system may be several journal entries behind those journal entries that are known to the source system. Using asynchronous delivery forces you to accept the risk that recovery can lose a number of journal entries given a failure on the source system. The only advantage of this approach is that it may have minimal impact to the journaling throughput on the source system when compared to synchronous delivery mode.

The difference between what journal entries exist in the remote journal on the target system from those residing in the journal on the source system is known as journal entry latency. Latency appears less on V4R2 systems than for most hot-site backup environments on systems prior to V4R2 (especially if OptiConnect/400 bus transport is used due to the high-speed connection). However, a synchronous solution that assures no latency, and no lost database updates, remains the preferred approach.

Note: Regarding journal entry latency for low-volume journal activity, it is possible that asynchronously maintained remote journals using a bus transport can be nearly as current as synchronously maintained remote journals. They can also cause less overall system overhead. However, for OptiConnected machines, this difference is barely measurable.

17.5 Performance Considerations for Remote Journal Implementation

There are several considerations when planning a remote journal environment. These include:

- Create separate ASPs for journal receivers.
- Direct the journal to remove internal entries from the journals.
- Direct the journal to employ a minimum fixed-length header.
- Make sure *BASE pool on the target system has enough storage assigned.

These are discussed further in the sections that follow.

17.5.1 Create Journal Receiver ASPs

To improve the performance of the journal function on both the source and target systems, create separate ASPs to be used exclusively for the receivers. When creating the ASPs, configure them with the fastest disk arms and as many arms as possible (up to a maximum of ten arms). The more disk arms you have in the ASP, the more parallel writes you have in journal receivers and the better your performance is.

Recommendation

Make sure that the ASP housing the corresponding target journal receivers has at least as many arms as the matching source journal receiver ASP.

Employ write cache to further enhance journal entry deposit performance.

Note: In the *AS/400 Advanced Series System Handbook*, GA19-5486, refer to the chapter on “Magnetic Media Controllers” to find a list of adapters that provide write caching, and to determine disk drive speeds.

17.5.2 Remove Internal Journal Entries

The receiver size options (RCVSIZOPT) parameter on the Change Journal (CHGJRN) command has a Remove Internal Entry (*RMVINTENT) option. Using this option inhibits sending hidden access path page images and implicit “before” images of database records to the target system. It also reduces unnecessary communications traffic. Because the access-path page images are 8K in size, using *RMVINTENT decreases the amount of data transported for more efficient use of the transport resource. *RMVINTENT also reduces the amount of storage required on both the source and target systems.

17.5.3 Reduce Length of Journal Entries

Another useful option of the RCVSIZOPT parameter is Minimum Fixed Length (*MINFIXLEN). There are approximately 46 fewer bytes sent to the source system’s disk for each journal entry deposit when using this option. These 46 bytes are not sent to the target system. The *MINFIXLEN value removes the job, program, and user profile information of each journal entry deposit.

Note: If you use *RCVSIZOPT or *MINFIXLEN with the RCVJRNE, RTVJRNE, and DSPJRNE commands, the job, program, and user profile information does not appear.

The CHGJRN parameters are shown in the following figure.

Change Journal (CHGJRN)		
Journal	JRN	
Library		*LIBL
Journal receiver:	JRNRCV	
Journal receiver		*SAME
Library		
Journal receiver		
Library		
Sequence option	SEQOPT	*CONT
Journal message queue	MSGQ	*SAME
Library		
Manage receivers	MNGRCV	*SAME
Delete receivers	DLTRCV	*SAME
Receiver size options	RCVSIZOPT	*RMVINTENT
		*MINFIXLEN
Journal state	JRNSTATE	*SAME
Type choices, press Enter.		

Figure 102. Receiver Size Options

Note: You can select one or both of the RCVSIZOPT option values on the Change Journal (CHGJRN) command. These options can also be selected when the journal is created using the Create Journal (CRTJRN) command.

17.5.4 Check *BASE Main Storage Pool Size on Target Machine

You should *always* have enough storage allocated to the *BASE storage pool. Implementing the remote journal function emphasizes this recommendation further, since many of the processes are performed at the MI level and use *BASE. When *BASE is sized large enough, machine functions are handled more efficiently. It is estimated that at least 10MB of main storage are required for the remote journal catch-up phase during an unplanned outage.

17.6 Performance Considerations for Running Remote Journal

One of the major performance enhancements of the remote journal function is the reduction in CPU overhead created by driving the transport mechanism functions into the microcode. The effect on the overall system performance depends on a number of factors. These include:

- **The transport method involved**

Use either an OptiConnect for OS/400 bus transport or a TCP/IP or SNA communications transport when using synchronous or asynchronous delivery. Use a non-OptiConnect transport when replication is over a long distance. When the delivery mode is synchronous, you need to weigh a slight increase in the response time against the communication overhead of the transport method. The most important factors regarding a communication transport method are the overall rated speed of the communications resource and any existing traffic already using the communications resource.

- **Number of remote journals that are maintained**

The impact to the job performing the database update increases by an equal factor for each (broadcasted) remote journal added. For example, three synchronously maintained remote journals, the impact to the job is three times that of one synchronously maintained journal.

By contrast, the impact is significantly less on the job performing the journal entry deposit for an asynchronously maintained remote journal.

Recommendation

Maintain only one synchronous remote journal for a given local journal since the application cannot continue until the journal entry is replicated to each remote journal.

- **Arrival rate of journal entries deposited on local system**

Journaling throughput for asynchronous and synchronous delivery is affected by the rate of journal entries deposited on the local system. Journal maintenance can fall further behind particularly for asynchronous journal transport mode.

- **Batch versus interactive initiation of journal maintenance**

In general, higher remote journal throughput can be maintained when many interactive jobs generate the journal throughput rather than a single-threaded batch job. This also requires less journal and remote journal overhead.

- **CPU usage on the source system**

The greater the CPU usage on the source system, the greater the potential of affecting journaling throughput. In applications constructed with heavy database update and insert activity, the extra CPU overhead attributed to remote journal on the source machine can range up to five or seven percent. It is even higher for synchronous delivery. With asynchronous delivery, high CPU levels can cause additional latency for updates.

- **CPU usage on the target system**

High CPU usage on the target system affects synchronous and asynchronous delivery modes the same way as described for the source system.

- **Sending task priority**

The higher the sending task priority value, the smaller the effect is that the remote journal function has on both the source and target systems. One exposure is that the target system can lag behind the source system. The sending task priority value can only be changed when using asynchronous delivery mode.

17.7 Remote Journal APIs

For V4R2 systems, there are five Application Program Interfaces (APIs) supplied.

- QjoAddRemoteJournal
- QjoChangeJournalState
- QjoRemoveRemoteJournal
- QjoRetrieveJournalInformation
- QjoRtvJrnReceiverInformation

When choosing to develop your own data replication applications, use these APIs to perform remote journal functions. Sample programs in C and RPG are found in the next section.

Note: Customers using a high availability business partner product do not need to implement code with these APIs. The business partner software handles these functions for you.

For additional information on these APIs, refer to the *System API Reference*, SC41-5801-01.

17.8 Remote Journal Coding Examples

Included in the *V4R2 Backup and Recovery* manual, SC41-5304, is a chapter on the remote journal function. There are scenarios about using the remote journal function for both data replication and hot-site backup environments, as well as remote journal recovery.

As you consider implementing remote journals in your availability and recovery plan, review the *Backup and Recovery* manual, SC41-5304, before implementing any of the examples described here. Remember, if you elect to implement DataPropagator/400, OMS, Mimix, or other high availability solutions, you *do not* need to write programming code to benefit from the remote journal function.

This section provides two examples for implementing remote journals—one in C and one in RPG. The source is provided on a “as is” and “best effort” basis. Thoroughly test the compiled programs in your operating environment.

17.8.1 Implementing Remote Journals with C

The example in this section lists the source statements for command, CL, and C programs to implement remote journal control interfaces using a command. Each program is self-documented with comments.

Note: A Qsysinc implementation is preferred to the example provided here. Using Qsysinc reduces the need for declares and provides more function with less coding.

The source is listed for five remote journal functions:

- Create a remote journal
- Start remote journal
- Activate remote journal
- De-activate remote journal
- Remove remote journal

Any handling of error messages is left to the user. “As is errors” perform a check function. The “As is errors” are written to the job log for the user to review and act upon.

Remember, you need the appropriate string and other header files in your implementation of this C program.

```

/* START REMOTE JRN CMD SRC */
CMD PROMPT(' CREATE REMOTE JRN')
    PARM          KWD(JRN) TYPE(*CHAR) LEN(10) MIN(1) +
                  PROMPT(' LOCAL JOURNAL NAME')
    PARM          KWD(LIB) TYPE(*CHAR) LEN(10) MIN(1) +
                  PROMPT(' JOURNAL LIB LOCAL+REM')
    PARM          KWD(RDB) TYPE(*CHAR) LEN(18) MIN(1) +
                  PROMPT(' ADDRDBDIRE TO USE')

/* START REMOTE JRN CL SRC */
PGM PARM(&JOURNAL &LIB &RDB)
DCL &JOURNAL TYPE(*CHAR) LEN(10)
DCL &LIB TYPE(*CHAR) LEN(10)
DCL &RDB TYPE(*CHAR) LEN(18)
DCL &JRNSTR TYPE(*CHAR) LEN(20)

CHGVAR &JRNSTR (&JOURNAL *CAT &LIB)
CALL RMJON2 (&JRNSTR &RDB)
ENDPGM

/* START REMOTE JRN C SRC PROGRAM RMJON2 */
/* To simplify the coding, the following assumptions are made: */
/* 1) The remote journal is created in the same named library */
/* as that of the source journal */
/* 2) The request variable structure, which is part of */
/* the Omissible Parameter Group, is not provided. Thus the */
/* remote journal receiver is created in the same named */
/* library as the source journal receiver on the source */
/* system and the remote journal type defaults to *TYPE2. */
/* Input parameters: */
/* argv[1] - Qualified Journal Name */
/* argv[2] - Relational DB Entry */

int main(int argc, char *argv[])
{
    char * reqval; /* Request variable */
    long * lnrq; /* Length of request variable */
    char * fmtname; /* Format name of request variable */
    long * ioerr; /* Error code */
    reqval = NULL; /* Since request variable is not */
    lnrq = NULL; /* set the pointers to NULL */
    fmtname = NULL;
    ioerr = NULL;

    QjoAddRemoteJournal(argv[1],argv[2],reqval,lnrq,fmtname,ioerr);
}

```

```

/* ACTIVATE REMOTE JRN COMMAND SOURCE */
CMD PROMPT(' ENGAGE REMOTE JRN')
PARM KWD(JRN) TYPE(*CHAR) LEN(10) MIN(1) +
PROMPT(' LOCAL JOURNAL NAME')
    PARM      KWD(LIB) TYPE(*CHAR) LEN(10) MIN(1) +
              PROMPT(' JOURNAL LIB SOURCE')
    PARM      KWD(RDB) TYPE(*CHAR) LEN(18) MIN(1) +
              PROMPT(' ADDRBDIRE TO USE')
    PARM      KWD(RMTJRN) TYPE(*CHAR) LEN(10) MIN(1) +
              PROMPT(' REMOTE JOURNAL NAME ')
    PARM      KWD(RMTLIB) TYPE(*CHAR) LEN(10) MIN(1) +
              PROMPT(' REMOTE JOURNAL LIB ')
    PARM      KWD(RMTRCV) TYPE(*CHAR) LEN(10) +
              DFT(*ATTACHED) MIN(0) PROMPT(' REMOTE +
              JOURNAL RCVR ')

/* ACTIVATE CL PGM SOURCE */
PGM PARM(&JOURNAL &LIB &RDB &RMTJRN &RMTLIB &RMTRCV)
DCL &JOURNAL TYPE(*CHAR) LEN(10)
DCL &LIB TYPE(*CHAR) LEN(10)
DCL &RDB TYPE(*CHAR) LEN(18)
DCL &RMTJRN TYPE(*CHAR) LEN(10)
DCL &RMTLIB TYPE(*CHAR) LEN(10)
DCL &RMTRCV TYPE(*CHAR) LEN(10)
DCL &JRNSTR TYPE(*CHAR) LEN(20)
DCL &RMTSTR TYPE(*CHAR) LEN(58)

CHGVAR &JRNSTR (&JOURNAL *CAT &LIB)
CHGVAR &RMTSTR (&RDB *CAT &RMTJRN *CAT &RMTLIB *CAT &RMTRCV *CAT &RMTLIB)
CALL RMJACT (&JRNSTR &RMTSTR)
ENDPGM

/* ACTIVATE C PGM SOURCE */
/* Input parameters: */
/* argv[1] - Qualified Journal Name */
/* argv[2] - Request variable */

int main(int argc, char *argv[])
{
    char fmt[8]; /* Format name of request variable */
    long fmtlen;

    char * recval; /* Receiver variable */
    long * lnrcv; /* Length of receiver variable */
    char * fmtname; /* Format name of request variable */
    long * ioerr; /* Error code */

    strcpy(fmt,"CJST0400");
    fmtname = fmt;
    fmtlen = 58;
    lnrcv = &fmtlen;
    lnrcv=NULL;
    recval = NULL;
    ioerr = NULL;

    QjoChangeJournalState(argv[1],argv[2],lnrcv,fmtname,recval,lnrcv,ioerr);
}

```

```

/* DEACTIVATE COMMAND SOURCE */
CMD PROMPT('DISENGAGE REMOTE JRN')
PARM KWD(JRN) TYPE(*CHAR) LEN(10) MIN(1) +
PROMPT('LOCAL JOURNAL NAME')
    PARM      KWD(LIB) TYPE(*CHAR) LEN(10) MIN(1) +
              PROMPT('JOURNAL LIB LOCAL+REM')
    PARM      KWD(RDB) TYPE(*CHAR) LEN(18) MIN(1) +
              PROMPT('ADDRDBDIRE TO USE')
    PARM      KWD(TYPE) TYPE(*CHAR) LEN(1) RSTD(*YES) +
              DFT('0') VALUES('0' '1') MIN(0) PROMPT('TYPE +
              0=CTL 1=IMMED')

/* CL PGM SOURCE */
PGM PARM(&JOURNAL &LIB &RDB &TYPE)
DCL &JOURNAL TYPE(*CHAR) LEN(10)
DCL &LIB TYPE(*CHAR) LEN(10)
DCL &RDB TYPE(*CHAR) LEN(18)
DCL &TYPE TYPE(*CHAR) LEN(1)
DCL &RMTLIB TYPE(*CHAR) LEN(10)
DCL &RMTRCV TYPE(*CHAR) LEN(10)
DCL &JRNSTR TYPE(*CHAR) LEN(20)
DCL &RMTSTR TYPE(*CHAR) LEN(39)

CHGVAR &JRNSTR (&JOURNAL *CAT &LIB)
CHGVAR &RMTSTR (&RDB *CAT &JOURNAL *CAT &LIB *CAT &TYPE)
CALL RMJDEACT (&JRNSTR &RMTSTR)
ENDPGM

/* C PGM SOURCE */
/* Input parameters: */
/* argv[1] - Qualified Journal Name */
/* argv[2] - Request variable */

int main(int argc, char *argv[])
{
    char fmt[8];
    long fmtlen;

    char * recval;      /* Receiver variable */
    long * lnrq;        /* Length of request variable */
    long * lnrcv;       /* Length of receiver variable */
    char * fmtname;     /* Format name of request variable */
    long * ioerr;       /* Error code */

    strcpy(fmt,"CJST0300");
    fmtname = fmt;
    fmtlen = 39;
    lnrq = &fmtlen;
    lnrcv=NULL;
    recval = NULL;
    ioerr = NULL;

    QjoChangeJournalState(argv[1],argv[2],lnrq,fmtname,recval,lnrcv,ioerr);
}

```

```

/* REMOVE JRN CMD SRC */
CMD PROMPT(' REMOVE REMOTE JRN')
  PARM KWD(JRN) TYPE(*CHAR) LEN(10) MIN(1) +
  PROMPT(' LOCAL JOURNAL NAME')
      PARM      KWD(LIB) TYPE(*CHAR) LEN(10) MIN(1) +
          PROMPT(' JOURNAL LIB LOCAL+REM')
      PARM      KWD(RDB) TYPE(*CHAR) LEN(18) MIN(1) +
          PROMPT(' ADDRDBDIRE TO USE')

/* END REMOTE JRN CL SRC */
PGM PARM(&JOURNAL &LIB &RDB)
DCL &JOURNAL TYPE(*CHAR) LEN(10)
DCL &LIB TYPE(*CHAR) LEN(10)
DCL &RDB TYPE(*CHAR) LEN(18)
DCL &JRNSTR TYPE(*CHAR) LEN(20)

CHGVAR &JRNSTR (&JOURNAL *CAT &LIB)
CALL RMJOFF (&JRNSTR &RDB)
ENDPGM
/* END REMOTE JRN C SRC */
/* Input parameters: */
/* argv[1] - Qualified Journal Name */
/* argv[2] - Relational DB Enrty */

int main(int argc, char *argv[])
{
    char * reqval; /* Request variable */
    long * lnrq; /* Length of request variable */
    char * fmtname; /* Format name of request variable */
    long * ioerr; /* Error code */
    reqval = NULL; /* Since request variable is not */
    lnrq = NULL; /* set the pointers to NULL */
    fmtname = NULL;
    ioerr = NULL;

    QjoRemoveRemoteJournal (argv[1],argv[2],reqval,lnrq,fmtname,ioerr);
}

```

17.8.2 Implementing Remote Journals with RPG

The example in this section lists the source specifications for three command programs and three RPG programs. Each program is self-documented with comments. The three command and program names are:

- ADDRMTJRN—to add a remote journal
- CHGJRNSTT—to change a remote journal state
- RMRMTJRN—to remove a remote journal

Error handling is left to the user.

Please note the command and program names are user created. To ensure your application uses the commands you create when executed, we recommend that you qualify the command with the associated library name. If IBM or vendors providing application interface or support for remote journals choose the same names as selected for this example, you can be sure that your commands and programs are used. This remains true only if the program name is qualified with the associated library name. As usual, we also recommend that you store the commands and programs in a library that will not be replaced during a release upgrade.

SOURCE FILE RMTJRN/QCMD SRC

MEMBER ADDRMTJRN

SEQNBR* 1 2 3 4 5 6 7 8 9 0

```

100      CMD      PROMPT('Add Remote Journal')
200      PARM     KWD(JRNNAME) TYPE(QJRNAME) MIN(1) +
300              PROMPT('Journal Name')
400      PARM     KWD(RDBDIRE) TYPE(*CHAR) LEN(18) MIN(1) +
500              PROMPT('Remote RDB Directory Entry')
600      PARM     KWD(RMTJRNAME) TYPE(QRMTJRNNM) +
700              PMTCTL(*PMTRQS) PROMPT('Remote Journal Name')
800      PARM     KWD(RJRNLIB) TYPE(*CHAR) LEN(10) +
900              DFT(*SRCSYS) SPCVAL((*SRCSYS)) +
1000     PMTCTL(*PMTRQS) PROMPT('Remote Journal +
1100     Rcvr Library')
1200     PARM     KWD(RMTJRNTYP) TYPE(*CHAR) LEN(1) RSTD(*YES) +
1300     DFT(1) VALUES(1 2) PMTCTL(*PMTRQS) +
1400     PROMPT('Remote Journal Type')
1500     PARM     KWD(MSGQNAME) TYPE(QMSGQNAME) +
1600     PMTCTL(*PMTRQS) PROMPT('Message Queue Name')
1700     PARM     KWD(DLTRCV) TYPE(*CHAR) LEN(1) RSTD(*YES) +
1800     DFT(N) VALUES(Y N) PMTCTL(*PMTRQS) +
1900     PROMPT('Delete Receivers?')
2000     PARM     KWD(RMTJRNTXT) TYPE(*CHAR) LEN(50) +
2100     PMTCTL(*PMTRQS) PROMPT('Remote Journal Text')
2200     QJRNAME:  QUAL  TYPE(*NAME) LEN(10) MIN(1)
2300     QUAL      TYPE(*NAME) LEN(10) SPCVAL((*CURLIB) +
2400     (*LIBL)) MIN(1) PROMPT('Library')
2500     QRMTJRNNM:  QUAL  TYPE(*NAME) LEN(10) DFT(*JRN) SPCVAL((*JRN)) +
2600     MIN(0)
2700     QUAL      TYPE(*NAME) LEN(10) MIN(0) PROMPT('Library')
2800     QMSGQNAME:  QUAL  TYPE(*NAME) LEN(10) DFT(QSYSOPR)
2900     QUAL      TYPE(*NAME) LEN(10) DFT(QSYS) PROMPT('Library')
          * * * * * E N D   O F   S O U R C E   * * * * *
    
```

SOURCE FILE RMTJRN/QRPGLESRC

MEMBER ADDRMTJRN

SEQNBR*...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+... 8 ...+... 9 ...+... 0

```

100 H***** 03/15/98
200 H* This program accepts input from a CMD to Add a Remote Journal. 03/15/98
300 H* /COPY is used to reduce the size of the source. Appropriate 03/15/98
400 H* parts of the includes can also be copied directly to this source 03/15/98
500 H* member. 03/15/98
600 H***** 03/15/98
700 H* Compile options 03/15/98
800 HDFTACTGRP(*NO) ACTGRP(*CALLER) 03/14/98
900 D***** 03/15/98
1000 D* Includes for Journal APIs 03/15/98
1100 D* 03/15/98
1200 D/COPY QSYSINC/QRPGLESRC,QJOURNAL 03/15/98
1300 D***** 03/15/98
1400 D* 03/15/98
1500 D***** 03/15/98
1600 D* Error code parameter include. Bytes available is set to 0 in *INZSR 03/15/98
1700 D* which causes program to function check if any errors are 03/15/98
1800 D* encountered. Error handling can be added to this program by 03/15/98
1900 D* changing the bytes available parameter to 16, or greater than 16. 03/26/98
2000 D* Changing this value without adding any error handling to the 03/15/98
2100 D* program will cause the program to appear to complete normally 03/15/98
2200 D* even when an error has occurred, other than messages logged to 03/15/98
2300 D* the joblog. 03/15/98
2400 D***** 03/15/98
2500 D* 03/15/98
2600 D***** 03/15/98
2700 D* Standalone fields 03/15/98
2800 D* 03/15/98
2900 D/COPY QSYSINC/QRPGLESRC,QUSEC 03/15/98
3000 D REQVARLEN S 9B 0 03/15/98
3100 D* Length of request variable 03/26/98
3200 D REQFMNAM S 8A 03/15/98
3300 D* Request format name 03/26/98
3400 D***** 03/15/98
3500 C***** 03/26/98
3600 C* Main Line 03/26/98
3700 C EXSR ADDJRN 03/15/98
3800 C MOVE *ON *INLR 03/15/98
3900 C* End of Main Line 03/26/98
4000 C***** 03/26/98
4100 C***** 03/26/98
4200 C* Program Initialization Subroutine 03/26/98
4300 C* 03/26/98
4400 C *INZSR BEGSR 03/15/98
4500 C *ENTRY PLIST 03/15/98
4600 C PARM JRNNAME 20 03/19/98
4700 C* Journal Name 03/26/98
4800 C PARM RDBDIRE 18 03/15/98
4900 C* Remote Relational DB 03/26/98
5000 C* Directory Entry 03/26/98
5100 C PARM RMTJRNAME 20 03/19/98
5200 C* Remote Journal Name 03/26/98
5300 C PARM RJRNRLIB 10 03/15/98
5400 C* Remote Journal Receiver 03/26/98
5500 C* Library 03/26/98
5600 C PARM RMTJRNTYP 1 03/15/98
5700 C* Remote Journal Type 03/26/98
5800 C* 1 = *TYPE1 03/26/98
5900 C* 2 = *TYPE2 03/26/98
6000 C PARM MSGQNAME 20 03/19/98

```

```

6100 C*                               Message Queue Name           03/26/98
6200 C                               PARM          DLTRCV          1           03/15/98
6300 C*                               Delete Receivers           03/26/98
6400 C*                               N = 0/Don't Delete         03/26/98
6500 C*                               Y = 1/Delete              03/26/98
6600 C                               PARM          RMTJRNTXT          50           03/15/98
6700 C*                               Remote Journal Text       03/26/98
6800 C                               EVAL          QUSBPRV = 0           03/15/98
6900 C*                               See comments in D specs for 03/26/98
7000 C*                               error code parameter.    03/26/98
7100 C*                               03/26/98
7200 C                               ENDSR                          03/15/98
7300 C*****                          03/26/98
7400 C* Add Remote Journal Subroutine 03/26/98
7500 C*                               03/26/98
7600 C   ADDJRN          BEGSR                          03/15/98
7700 C                               EVAL          REQVARLEN = %SIZE(QJOJ0100) 03/15/98
7800 C                               EVAL          REQFMTNAM = 'ADRJ0100' 03/15/98
7900 C   RMTJRNAME      IFNE          '*JRN'           03/26/98
8000 C                               EVAL          QJOQRJN  = RMTJRNAME 03/19/98
8100 C*                               Special value *JRN indicates remote 03/26/98
8200 C*                               journal name is the same as the local 03/26/98
8300 C*                               journal name.  If this parameter = *JRN 03/26/98
8400 C*                               leave QJOQRJN blank.  Otherwise set to 03/26/98
8500 C*                               passed name.                    03/26/98
8600 C                               ENDF                          03/26/98
8700 C   RJRNRLIB       IFNE          '*SRCSYS'         03/26/98
8800 C                               EVAL          QJORJRLN = RJRNRLIB 03/15/98
8900 C*                               Special value *SRCSYS indicates remote 03/26/98
9000 C*                               journal receiver library is the same as 03/26/98
9100 C*                               the library on the source system.  If this 03/26/98
9200 C*                               parameter = *SRCSYS, leave QJORJRLN 03/26/98
9300 C*                               blank.  Otherwise set to passed name. 03/26/98
9400 C                               ENDF                          03/26/98
9500 C                               EVAL          QJORJT   = RMTJRNTYP 03/15/98
9600 C                               EVAL          QJOQJMQ  = MSGQNAME 03/19/98
9700 C   DLTRCV        IFEQ          'Y'              03/19/98
9800 C*                               Delete Receivers = Yes         03/26/98
9900 C                               EVAL          QJODR    = '1'         03/19/98
10000 C                               ENDF                          03/19/98
10100 C   DLTRCV        IFEQ          'N'              03/19/98
10200 C                               EVAL          QJODR    = '0'         03/19/98
10300 C*                               Delete Receivers = No         03/26/98
10400 C                               ENDF                          03/19/98
10500 C                               EVAL          QJOTEXT  = RMTJRNTXT 03/15/98
10600 C                               CALLB         QJOARJ           02/02/98
10700 C                               PARM          JRNNAME           03/19/98
10800 C                               PARM          RBBDIR           03/15/98
10900 C                               PARM          QJOJ0100         03/15/98
11000 C                               PARM          REQVARLEN         03/15/98
11100 C                               PARM          REQFMTNAM         03/15/98
11200 C                               PARM          QUSEC           03/14/98
11300 C                               ENDSR                          03/15/98
11400 C*****                          03/26/98
      * * * * E N D O F S O U R C E * * * *

```


SOURCE FILE RMTJRN/QCMD SRC

MEMBER CHGJRNSTT

SEQNBR*...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+... 8 ...+... 9 ...+... 0

```

100      CMD      PROMPT(' Change Journal State')                                04/10/98
200      PARM     KWD(STATUS) TYPE(*CHAR) LEN(1) RSTD(*YES) +                    03/19/98
300          VALUES(A I) MIN(1) PROMPT(' Activate +                               04/09/98
400          or Inactivate?')                                                    03/17/98
500      PARM     KWD(LCLRMT) TYPE(*CHAR) LEN(1) RSTD(*YES) +                    03/17/98
600          VALUES(L R) MIN(1) PROMPT(' Local Or +                               04/09/98
700          Remote?')                                                            03/17/98
800      PARM     KWD(JRNNAME) TYPE(QJRNAME) MIN(1) +                             03/19/98
900          PROMPT(' Journal Name')                                              03/15/98
1000     PARM     KWD(RBBDIRE) TYPE(*CHAR) LEN(18) MIN(0) +                       03/26/98
1100          PMTCTL(REMOTE) PROMPT(' Remote RDB +                               03/26/98
1200          Directory Entry')                                                    03/26/98
1300     PARM     KWD(RJRNAME) TYPE(QJRNAME) PMTCTL(REMOTE) +                     03/26/98
1400          PROMPT(' Remote Journal Name')                                       03/26/98
1500     PARM     KWD(STRJRNRCV) TYPE(QSTRJRNRCV) +                               04/09/98
1600          DFT(*ATTACHED) SNGVAL((*ATTACHED) +                               04/09/98
1700          (*SRCSYS)) PMTCTL(RMTACT) +                                         04/09/98
1800          PROMPT(' Starting Journal Receiver Name')                           04/09/98
1900     PARM     KWD(PRFINACTYP) TYPE(*CHAR) LEN(1) +                           04/09/98
2000          RSTD(*YES) DFT(C) VALUES(C I) +                                     04/09/98
2100          PMTCTL(RMTINACT) PROMPT(' Preferred +                               04/09/98
2200          Inactive Type')                                                       04/09/98
2300     PARM     KWD(SYNASY) TYPE(*CHAR) LEN(1) RSTD(*YES) +                     04/09/98
2400          DFT(A) VALUES(S A) PMTCTL(RMTACT) +                               04/09/98
2500          PROMPT(' Synch or Async Maintained?')                               04/09/98
2600     PARM     KWD(SNDTSKPRI) TYPE(*INT4) DFT(0) +                             04/09/98
2700          PMTCTL(ASYNC) PROMPT(' Sending Task Priority')                       04/09/98
2800     QJRNAME:  QUAL  TYPE(*NAME) LEN(10) MIN(1)                               03/26/98
2900          TYPE(*NAME) LEN(10) SPCVAL((*LIBL) +                               03/26/98
3000          (*CURLIB)) MIN(1) PROMPT(' Library')                                 03/26/98
3100     QJRNAME:  QUAL  TYPE(*NAME) LEN(10) DFT(*JRN) SPCVAL((*JRN)) +          03/26/98
3200          MIN(0)                                                                03/26/98
3300          TYPE(*NAME) LEN(10) PROMPT(' Library')                               03/26/98
3400     QSTRJRNRCV:  QUAL  TYPE(*NAME) LEN(10) MIN(1)                           04/08/98
3500          TYPE(*NAME) LEN(10) SPCVAL((*LIBL) +                               03/26/98
3600          (*CURLIB)) PROMPT(' Library')                                       03/26/98
3700     REMOTE:   PMTCTL  CTL(LCLRMT) COND((*EQ R))                               03/26/98
3800     RMTACT:   PMTCTL  CTL(LCLRMT) COND((*EQ R))                               03/26/98
3900          PMTCTL  CTL(STATUS) COND((*EQ A))                                   03/26/98
4000     RMTINACT: PMTCTL  CTL(LCLRMT) COND((*EQ R))                               03/26/98
4100          PMTCTL  CTL(STATUS) COND((*EQ I))                                   03/26/98
4200     ASYNC:    PMTCTL  CTL(SYNASY) COND((*EQ A))                               03/26/98
4300          PMTCTL  CTL(STATUS) COND((*EQ A))                                   04/09/98
4400          PMTCTL  CTL(LCLRMT) COND((*EQ R))                                   04/10/98
          * * * * E N D   O F   S O U R C E   * * * *

```

SOURCE FILE RMTJRN/QRPGLESRC

MEMBER CHGJRNSTT

```

SEQNBR*...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+... 8 ...+... 9 ...+... 0
100 H*****
200 H* This program accepts input from a CMD to Change Journal State. 03/15/98
300 H* /COPY is used to reduce the size of the source. Appropriate 04/10/98
400 H* parts of the includes can also be copied directly to this source 03/15/98
500 H* member. 03/15/98
600 H* 03/19/98
700 H* This program only supports changing the journal state of local 03/19/98
800 H* journals on the source system and remote journals from the 03/19/98
900 H* source system. Inactivation of a remote journal from the 03/19/98
1000 H* target system is not coded in this example. This is primarily 03/19/98
1100 H* due to the fact that this would make the command more confusing 03/19/98
1200 H* if this support were added here. Format CJST0200 is required 03/19/98
1300 H* to inactivate a remote journal from a target system. 03/19/98
1400 H* 03/19/98
1500 H* CHGJRNSTT command uses selective prompting so that only applicable 04/10/98
1600 H* parameters are shown. 03/26/98
1700 H* 03/26/98
1800 H***** 03/15/98
1900 H* Compile options 03/15/98
2000 H* 04/09/98
2100 HDFTACTGRP(*NO) ACTGRP(*CALLER) 03/14/98
2200 D***** 03/15/98
2300 D* Includes for Journal APIs 03/15/98
2400 D* 03/15/98
2500 D/COPY QSYSINC/QRPGLESRC,QJOURNAL 03/15/98
2600 D***** 03/15/98
2700 D* 03/15/98
2800 D***** 03/15/98
2900 D* Error code parameter include. Bytes available is set to 0 in *INZSR 03/15/98
3000 D* which causes program to function check if any errors are 03/15/98
3100 D* encountered. Error handling can be added to this program by 03/15/98
3200 D* changing the bytes available parameter to 16, or greater than 16. 03/26/98
3300 D* Changing this value without adding any error handling to the 03/15/98
3400 D* program will cause the program to appear to complete normally 03/15/98
3500 D* even when an error has occurred, other than messages logged to 03/15/98
3600 D* the joblog. 03/15/98
3700 D/COPY QSYSINC/QRPGLESRC,QUSEC 03/18/98
3800 D***** 03/15/98
3900 D* 03/15/98
4000 D***** 04/08/98
4100 D* Includes for Retrieve Object Description API 04/08/98
4200 D* Used when *LIBL or *CURLIB is used for Journal Name and *JRN 04/08/98
4300 D* is used for Remote Journal Name 04/08/98
4400 D* 04/08/98
4500 D/COPY QSYSINC/QRPGLESRC,QUSROBJD 04/08/98
4600 D***** 04/08/98
4700 D* 04/08/98
4800 D***** 03/15/98
4900 D* Standalone fields 03/15/98
5000 D* 03/15/98
5100 D REQVARLEN S 9B 0 03/15/98
5200 D* Length of request variable 03/19/98
5300 D RCVVARLEN S 9B 0 03/18/98
5400 D* Length of receiver variable 03/19/98
5500 D REQFMTNAM S 8A 03/15/98
5600 D* Request format name 03/19/98
5700 D SNDTSKPRI S 9B 0 03/18/98
5800 D* Sending task priority 03/19/98
5900 DJRNOBJTYP S 10A INZ('*JRN') 04/08/98
6000 D* Object type of journal 04/08/98
6100 D***** 03/15/98

```



```

11200 C*****
11300 C* Check to see if *JRN specified for remote journal. If so, check to
11400 C* see if *LIBL or *CURLIB is specified for local journal name. If
11500 C* so, use Retrieve Object Description API to find local journal and
11600 C* change *LIBL or *CURLIB to library name where the journal was
11700 C* found. This is necessary because the API requires a library
11800 C* name for the remote journal and does not support *LIBL or *CURLIB
11900 C*
12000 C      RMTJRNNAME      IFEQ      '*JRN'
12100 C      '*LIBL'        SCAN      JRNNAME      SCANRESULT      2 0 90
12200 C      *IN90          IFNE      *ON
12300 C      '*CURLIB'     SCAN      JRNNAME      SCANRESULT      90
12400 C
12500 C      *IN90          IFEQ      *ON
12600 C
12700 C      EXSR          RTVLIB
12800 C
12900 C      ENDIF
13000 C*****
13100 C* End program initialization subroutine
13200 C      ENDSR
13300 C*****
13400 C* Local Journal Subroutine
13500 C
13600 C      LCLJRN        BEGSR
13700 C      EVAL          REQVARLEN = %SIZE(QJ00100R00)
13800 C      EVAL          REQFMNTNAM = 'CJST0100'
13900 C      EVAL          RCVVARLEN = 0
14000 C      STATUS        IFEQ          'I'
14100 C*              Status = Inactive
14200 C      EVAL          QJONJS = '0'
14300 C      ENDIF
14400 C      STATUS        IFEQ          'A'
14500 C*              Status = Active
14600 C      EVAL          QJONJS = '1'
14700 C      ENDIF
14800 C      CALLB        QJOCJS
14900 C      PARM          JRNNAME
15000 C      PARM          QJ00100R00
15100 C      PARM          REQVARLEN
15200 C      PARM          REQFMNTNAM
15300 C      PARM          QJ00100R00
15400 C      PARM          RCVVARLEN
15500 C      PARM          QUSEC
15600 C      ENDSR
15700 C*****
15800 C* Remote Journal Subroutine
15900 C
16000 C      RMTJRN        BEGSR
16100 C      STATUS        CASEQ          'I'          RMTINACT
16200 C*              Status = Inactive
16300 C      STATUS        CASEQ          'A'          ACTRMTJRN
16400 C*              Status = Active
16500 C      ENDCS
16600 C      ENDSR
16700 C*****
16800 C* Inactivate Remote Journal Subroutine
16900 C*
17000 C* This is the only call to the QjoChangeJournalState API that
17100 C* provides a receiver variable. The use of this receiver is
17200 C* not demonstrated in this example, but would be used to
17300 C* record information about the inactivation, such as receiver
17400 C* name and sequence number which could be used when restarting
17500 C* remote journaling.
17600 C
17700 C      RMTINACT      BEGSR
17800 C      EVAL          REQVARLEN = %SIZE(QJ00300R)
17900 C      EVAL          REQFMNTNAM = 'CJST0300'
18000 C      EVAL          RCVVARLEN = %SIZE(QJ00300R00)
18100 C      EVAL          QJORBDE = RDBDIRE

```

```

18100 C RMTJRNAME IFEQ '*JRN' 03/26/98
18200 C EVAL QJORJ = JRNNAME 03/26/98
18300 C* If RMTJRNAME = *JRN, set QJORJ = local journal name. 03/26/98
18400 C* Otherwise, set to name passed. 03/26/98
18500 C ELSE 03/26/98
18600 C EVAL QJORJ = RMTJRNAME 03/19/98
18700 C ENDIF 03/26/98
18800 C PRFINACTYP IFEQ 'C' 03/18/98
18900 C EVAL QJOPIT = '0' 03/18/98
19000 C* Preferred Inactive Type Controlled 03/19/98
19100 C ENDIF 03/18/98
19200 C PRFINACTYP IFEQ 'I' 03/26/98
19300 C EVAL QJOPIT = '1' 03/18/98
19400 C* Preferred Inactive Type Immediate 03/19/98
19500 C ENDIF 03/18/98
19600 C CALLB QJOCJS 03/18/98
19700 C PARM JRNNAME 03/19/98
19800 C PARM QJ00300R 03/18/98
19900 C PARM REQVARLEN 03/18/98
20000 C PARM REQFMTNAM 03/18/98
20100 C PARM QJ00300R00 03/18/98
20200 C PARM RCVVARLEN 03/18/98
20300 C PARM QUSEC 03/18/98
20400 C ENDSR 03/18/98
20500 C***** 03/19/98
20600 C* Activate Remote Journal Subroutine 03/19/98
20700 C* 03/19/98
20800 C ACTRMTJRN BEGSR 03/19/98
20900 C SYNASY CASEQ 'S' ACTRMTSYN 03/18/98
21000 C* Synchronous Maintained Remote Journal 03/26/98
21100 C SYNASY CASEQ 'A' ACTRMTASY 03/18/98
21200 C* Asynchronous Maintained Remote Journal 03/26/98
21300 C ENDCS 03/18/98
21400 C ENDSR 03/18/98
21500 C***** 03/19/98
21600 C* Activate Remote Journal Synchronously Subroutine 03/19/98
21700 C* 03/19/98
21800 C ACTRMTSYN BEGSR 03/18/98
21900 C EVAL REQVARLEN = %SIZE(QJ00400R) 03/18/98
22000 C EVAL REQFMTNAM = 'CJST0400' 03/18/98
22100 C EVAL RCVVARLEN = 0 03/18/98
22200 C EVAL QJORBDE01 = RDBDIRE 03/18/98
22300 C RMTJRNAME IFEQ '*JRN' 03/26/98
22400 C EVAL QJORJ01 = JRNNAME 03/26/98
22500 C* If RMTJRNAME = *JRN, set QJORJ01 = local journal name. 03/26/98
22600 C* Otherwise, set to name passed. 03/26/98
22700 C ELSE 03/26/98
22800 C EVAL QJORJ01 = RMTJRNAME 03/26/98
22900 C ENDIF 03/26/98
23000 C EVAL QJOSJR = STRJRNRCV 03/19/98
23100 C CALLB QJOCJS 03/18/98
23200 C PARM JRNNAME 03/19/98
23300 C PARM QJ00400R 03/18/98
23400 C PARM REQVARLEN 03/18/98
23500 C PARM REQFMTNAM 03/18/98
23600 C PARM QJ00400R 03/18/98
23700 C PARM RCVVARLEN 03/18/98
23800 C PARM QUSEC 03/18/98
23900 C ENDSR 03/18/98

```

```

24000 C*****
24100 C* Activate Remote Journal Asynchronously Subroutine 03/19/98
24200 C* 03/19/98
24300 C ACTRMTASY BEGSR 03/18/98
24400 C EVAL REQVARLEN = %SIZE(QJ00500R) 03/18/98
24500 C EVAL REQFMTNAM = 'CJST0500' 03/18/98
24600 C EVAL RCVVARLEN = 0 03/18/98
24700 C EVAL QJORBDE02 = RDBDIRE 03/19/98
24800 C RMTJRNAME IFEQ '*JRN' 03/26/98
24900 C EVAL QJORJ02 = JRNNAME 03/26/98
25000 C* If RMTJRNNAME = *JRN, set QJORJ02 = local journal name. 03/26/98
25100 C* Otherwise, set to name passed. 03/26/98
25200 C ELSE 03/26/98
25300 C EVAL QJORJ02 = RMTJRNAME 03/26/98
25400 C ENDIF 03/26/98
25500 C EVAL QJOSJ00 = STRJRNRCV 03/19/98
25600 C EVAL QJOSTP = SNDSKPRI 03/18/98
25700 C CALLB QJOCJS 03/18/98
25800 C PARM JRNNAME 03/19/98
25900 C PARM QJ00500R 03/18/98
26000 C PARM REQVARLEN 03/18/98
26100 C PARM REQFMTNAM 03/18/98
26200 C PARM QJ00500R 03/18/98
26300 C PARM RCVVARLEN 03/18/98
26400 C PARM QUSEC 03/18/98
26500 C ENDSR 03/18/98
26600 C***** 04/08/98
26700 C* This subroutine is called when *LIBL or *CURLIB is specified for 04/09/98
26800 C* the journal name, and *JRN is specified for the remote journal 04/09/98
26900 C* name. The Retrieve Object Description API is used to find the 04/09/98
27000 C* local journal by the library list or current library, and then 04/09/98
27100 C* return the name of the library where the local journal was found. 04/09/98
27200 C* The returned library is moved to the last 10 positions of the 04/09/98
27300 C* JRNNAME field. This enables other subroutines that deal with 04/09/98
27400 C* *JRN to find a library name 04/09/98
27500 C* 04/09/98
27600 C RTVLIB BEGSR 04/08/98
27700 C EVAL RCVVARLEN = %SIZE(QUSD0100) 04/08/98
27800 C EVAL REQFMTNAM = 'OBJD0100' 04/08/98
27900 C CALL QUSROBJD 04/08/98
28000 C PARM QUSD0100 04/08/98
28100 C PARM RCVVARLEN 04/08/98
28200 C PARM REQFMTNAM 04/08/98
28300 C PARM JRNNAME 04/08/98
28400 C PARM JRNOBJTYP 04/08/98
28500 C PARM QUSEC 04/08/98
28600 C MOVE QUSRL01 JRNNAME 04/08/98
28700 C ENDSR 04/08/98
      * * * * E N D O F S O U R C E * * * *

```

```

5769PW1 V4R2M0 980228          SEU SOURCE LISTING          04/08/98 20:22:27          PAGE 1
SOURCE FILE . . . . . RMTJRN/QCMDSRC
MEMBER . . . . . RMVMTJRN
SEQNBR*...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+... 8 ...+... 9 ...+... 0
100      CMD      PROMPT('Remove Remote Journal')          04/08/98
200      PARM      KWD(JRNNAME) TYPE(QJRNAME) MIN(1) +      03/19/98
300                                PROMPT('Journal Name')    03/19/98
400      PARM      KWD(RBDDIRE) TYPE(*CHAR) LEN(18) MIN(1) + 03/15/98
500                                PROMPT('Remote RDB Directory Entry') 03/15/98
600      PARM      KWD(RMTJRNAME) TYPE(QRMTJRNM) +          03/19/98
700                                PMTCTL(*PMTRQS) PROMPT('Remote Journal Name') 03/19/98
800  QJRNAME:  QUAL  TYPE(*NAME) LEN(10) MIN(1)              03/19/98
900      QUAL      TYPE(*NAME) LEN(10) SPCVAL((*CURLIB) +    04/08/98
1000     (*LIBL)) MIN(1) PROMPT('Library')                  04/08/98
1100  QRMTJRNM:  QUAL  TYPE(*NAME) LEN(10) DFT(*JRN) SPCVAL((*JRN)) 03/26/98
1200      QUAL      TYPE(*NAME) LEN(10) PROMPT('Library')    03/26/98
          * * *   E N D   O F   S O U R C E   * * *

```

```

5769PW1 V4R2M0 980228          SEU SOURCE LISTING          04/08/98 20:27:22          PAGE 1
SOURCE FILE . . . . . RMTJRN/QRPGLESRC
MEMBER . . . . . RMVMTJRN
SEQNBR*...+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+... 8 ...+... 9 ...+... 0
100      H*****
200      H* This program accepts input from a CMD to Remove a Remote Journal.          03/15/98
300      H* /COPY is used to reduce the size of the source. Appropriate              03/15/98
400      H* parts of the includes can also be copied directly to this source          03/15/98
500      H* member.                                                                    03/26/98
600      H*                                                                              03/26/98
700      H* Removing a remote journal only removes the definition on the source        03/26/98
800      H* system. If the journal is to be deleted from the remote system,          03/26/98
900      H* it must be deleted by the user.                                           03/26/98
1000     H*****
1100     H* Compile options
1200     HDFTACTGRP(*NO) ACTGRP(*CALLER)
1300     D*****
1400     D* Includes for Journal APIs
1500     D*
1600     D/COPY QSYSINC/QRPGLESRC,QJOURNAL
1700     D*****
1800     D*
1900     D*****
2000     D* Error code parameter include. Bytes available is set to 0 in *INZSR
2100     D* which causes program to function check if any errors are
2200     D* encountered. Error handling can be added to this program by
2300     D* changing the bytes available parameter to 16, or greater than 16.
2400     D* Changing this value without adding any error handling to the
2500     D* program will cause the program to appear to complete normally
2600     D* even when an error has occurred, other than messages logged to
2700     D* the job log.
2800     D/COPY QSYSINC/QRPGLESRC,QUSEC
2900     D*****

```

```

3000 D* 03/15/98
3100 D***** 03/15/98
3200 D* Standalone fields 03/15/98
3300 D* 03/15/98
3400 D REQVARLEN S 9B 0 03/15/98
3500 D* Length of request variable 03/26/98
3600 D REQFMTNAM S 8A 03/15/98
3700 D* Request format name 03/26/98
3800 D***** 03/15/98
3900 C***** 03/26/98
4000 C* Main Line 03/26/98
4100 C EXSR RMVJRN 03/15/98
4200 C MOVE *ON *INLR 03/15/98
4300 C* End of Main Line 03/26/98
4400 C***** 03/26/98
4500 C***** 03/26/98
4600 C* Program Initialization Subroutine 03/26/98
4700 C* 03/26/98
4800 C *INZSR BEGSR 03/15/98
4900 C *ENTRY PLIST 03/15/98
5000 C PARM JRNNAME 20 03/19/98
5100 C* Journal Name 03/26/98
5200 C PARM RDBDIRE 18 03/15/98
5300 C* Remote Relational DB 03/26/98
5400 C* Directory Entry 03/26/98
5500 C PARM RMTJRNNAME 20 03/19/98
5600 C* Remote Journal Name 03/26/98
5700 C EVAL QUSBPRV = 0 03/15/98
5800 C* See comments in D specs for 03/26/98
5900 C* error code parameter. 03/26/98
6000 C* 03/26/98
6100 C ENDSR 03/15/98
6200 C***** 03/26/98
6300 C* Remove Remote Journal Subroutine 03/26/98
6400 C* 03/26/98
6500 C RMVJRN BEGSR 03/26/98
6600 C EVAL REQVARLEN = %SIZE(QJOJ010000) 03/15/98
6700 C EVAL REQFMTNAM = 'RMRJ0100' 03/15/98
6800 C RMTJRNNAME IFNE '*JRN' 04/09/98
6900 C EVAL QJOQRJN00 = RMTJRNNAME 04/09/98
7000 C* Special value *JRN indicates remote 04/09/98
7100 C* journal name is the same as the local 04/09/98
7200 C* journal name. If this parameter = *JRN 04/09/98
7300 C* leave QJOQRJN blank. Otherwise set to 04/09/98
7400 C* passed name. 04/09/98
7500 C ENDF 03/26/98
7600 C CALLB QJORRJ 03/15/98
7700 C PARM JRNNAME 03/19/98
7800 C PARM RDBDIRE 03/15/98
7900 C PARM QJOJ010000 03/15/98
8000 C PARM REQVARLEN 03/15/98
8100 C PARM REQFMTNAM 03/15/98
8200 C PARM QUSEC 03/14/98
8300 C ENDSR 03/15/98
8400 C***** 03/26/98
*****
* * * * * E N D O F S O U R C E * * * * *

```

17.9 More Information on the Remote Journal Function

For additional information on remote journals, refer to the V4R2 edition of:

- *Backup and Recovery Guide*, SC41-5304
- *DB2 for OS/400 Database Programming*, SC41-5701-01
- *DB2 for OS/400 Programming Book*, SC41-5611-01
- *DB2/400 Advanced Functions*, SG24-4249-01
- *System API Reference*, SC41-5801-01
- “What’s New With DB2” article in the March 1998 issue of the *AS/400 Magazine*.

Information on the *AS/400 Magazine* can be found at the URL:

<http://www.AS400.ibm.com/magazine>

Chapter 18. OptiConnect for OS/400

This section provides an overview of OptiConnect for OS/400 and the design of a high availability clustered solution from a hardware viewpoint. This chapter explains how an AS/400 clustered environment is linked and the various hardware options that it requires.

18.1 What OptiConnect for OS/400 Is

OptiConnect for OS/400 is an AS/400-to-AS/400 communication clustering solution. It combines unique OptiConnect fiber bus hardware and standard AS/400 bus hardware with unique software. This configuration allows for high-speed (1063 Mbps or 266 Mbps) communication using optical link AS/400 shared buses.

OptiConnect for OS/400 uses distributed data management (DDM) to allow applications on one AS/400 system to access databases located on other AS/400 systems. The AS/400 systems that contain the databases are the database servers, and the remote systems are considered the application clients or clients.

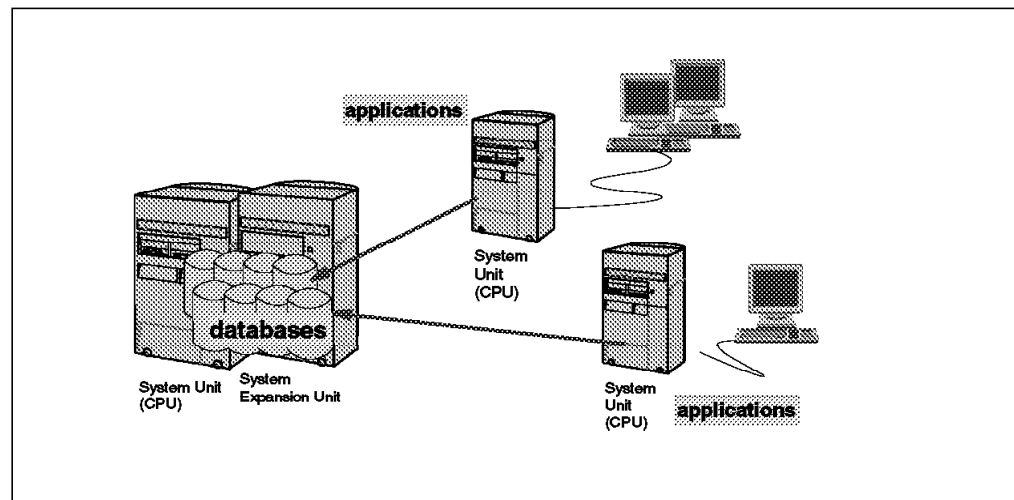


Figure 103. OptiConnect Database Server and Application Clients

In most cases, the hub also acts as the database server. Since all systems can communicate with each other (providing that the hub is active), any system can be the client. Some OptiConnect configurations have AS/400 systems that act simultaneously as a server and a client. However, any system can act as a database server.

The satellite system uses standard AS/400 fiber bus hardware, which may already exist in the installed system or need to be ordered.

OptiConnect for OS/400 supports the clustering of selected Reduced Instruction Set Computing (RISC) and Complex Instruction Set Computing (CISC) 9406 AS/400 processor models. A given cluster can include up to 32 AS/400 systems, with a mixture of both RISC and CISC architectures. Interoperability between OS/400 versions is maintained so each system can operate at different release levels. This assumes that all the systems are at V3R1 or later.

To determine the appropriate adapter for your AS/400 system, refer to Table 41 on page 359, which contains a list of applicable feature codes.

Note: OptiConnect for OS/400 is a communication vehicle. OptiConnect for OS/400 products provide AS/400 systems with *physical* links for a high availability clustered solution. OptiConnect for OS/400 components support the infrastructure for applications to conduct data exchanges over high speed connections.

OptiConnect for OS/400 does *not* offer high availability without applications that utilize the hardware links. OptiConnect for OS/400 can be the transport mechanism for in-house developed applications, business partner software, or remote journal support. See Appendix E, “High Availability Solutions” on page 391, to learn more about several high availability solutions highlighted in this redbook.

Note: For the purpose of our discussion, this chapter focuses on RISC systems only.

18.2 OptiConnect Hardware

From an AS/400 hardware point of view, an OptiConnect configuration or cluster is a collection of AS/400 systems connected with dedicated system buses using fiber optic cables. An OptiConnect cluster consists of at least one hub and one or more satellites. The AS/400 hub dedicates a system expansion unit to house the OptiConnect adapters. The shared bus allows all AS/400 systems to communicate with one another.

After the release of V4R2 and prior to the next release, you can order some optical link adapters as a standard hardware features. Therefore, you can order the hardware without needing special approval from the IBM laboratory. For a list of the required RPQ or feature code, refer to Table 41 on page 359.

OptiConnect clusters support greater flexibility for disaster recovery and high-availability solutions with the release of V4R2. The OptiConnect 2683 adapter was re-engineered to support a maximum optical link length of two kilometers at a speed of 266 Mbps. The prior limitation was 500 meters with a 2685 adapter at a 1063 Mbps speed.

The 2683 adapter also supports redundant fiber link solutions between the system unit and the adapter. This is increasingly important for availability as link distance increases.




Figure 104. Link Redundancy

For a detailed introduction to the terminology and hardware components of an OptiConnect high-availability solution, see Appendix D, “OptiConnect for OS/400 Terminology and Hardware Overview” on page 387.

18.3 OptiConnect for OS/400 Software Component Overview

OptiConnect for OS/400 software is comprised of two separate components:

- OptiConnect/400 for OS/400 – the base OptiConnect software functions
- OptiMover—a software component that provides an API interface to OptiConnect hardware

18.3.1 OptiConnect for OS/400

The OptiConnect for OS/400 software supports both remote journal and ObjectConnect for OS/400 over the OptiConnect hardware. The OptiConnect solution is required for the DDM interface to the optical link hardware. This hardware is a necessary component of the OptiConnect for OS/400 solution.

OptiConnect for OS/400 is a chargeable option of OS/400. One license is required for each AS/400 system in the cluster.

To confirm that OptiConnect for OS/400 is installed on your system, use the Display Software Resources (DSPSFWRSC) command. Check the resulting list for a resource ID of 57nn-SS1, a product option of 23, and a description of *OptiConnect*.

18.3.2 OptiMover for OS/400 (PRPQ P84291 Product Number 5799-FWQ)

The OptiMover software component provides an API interface to the OptiConnect hardware. The component allows you to tailor applications, and does not require OptiConnect for OS/400 software. Note that one OptiMover license is required for each AS/400 system.

OptiMover for OS/400 is a specialized version of OptiConnect for OS/400 that provides application APIs. Because OptiMover is essentially the same as the OS/400 OptiConnect feature, it uses the same combination of hardware and

software that allows you to connect multiple AS/400 systems on a shared optical bus.

With the OptiMover for OS/400 PRPQ, applications can use the high-speed optical bus to communicate among systems on the shared bus. This includes the object transfer functions in ObjectConnect. The high speed DDM capabilities of OptiConnect are not available, however. These APIs are available free of charge to OptiConnect for OS/400 customers by ordering a PTF that serves as an ordering vehicle.

OptiMover for OS/400 is available for those AS/400 RISC systems that require access to APIs to satisfy an application program prerequisite, but do not require full DDM support provided with OptiConnect. Use OS/400 option 22 (which supplies the ObjectConnect/400 commands) with OptiMover when you need an effective way to move objects between systems at optical bus speeds. Refer to Section 5.9, "ObjectConnect/400" on page 65 for information on ObjectConnect/400.

To confirm that OptiMover is installed on your system, use:

```
CHKPRDOPT PRDID(5799FWQ)
```

or

```
DSPPTF LICPGM(5799FWQ)
```

For more information, refer to *OptiMover for OS/400*, SC41-0626.

18.4 Bus Technology and the OptiConnect Environment

Figure 103 on page 351 shows how systems are connected together with fiber. What if the systems are connected using OptiConnect buses? How does the AS/400 system see the OptiConnect buses and how do these special OptiConnect cards work?

This section answers these questions and more. It assumes you are familiar with such terms as *hub*, *satellite*, *link*, *dual path*, and *path redundancy*. For a review of OptiConnect terminology, see Appendix D, "OptiConnect for OS/400 Terminology and Hardware Overview" on page 387.

A typical AS/400 bus always starts with an optical bus receiver card. Does it work if the bus receiver card is placed in the middle of the bus? In this case, the answer is yes. Each OptiConnect card is a special bus receiver card and acts as the primary bus receiver.

The attached system shows the system expansion tower as its own bus because the OptiConnect card reports in as a bus receiver. Each card in the hub represents a logical hardware resource starting with the hub's bus receiver card in slot 1 (logical slot 0). This information is shown using the hardware service manager option on the System Service Tools Logical Hardware Resources display as presented in Figure 105 on page 355.

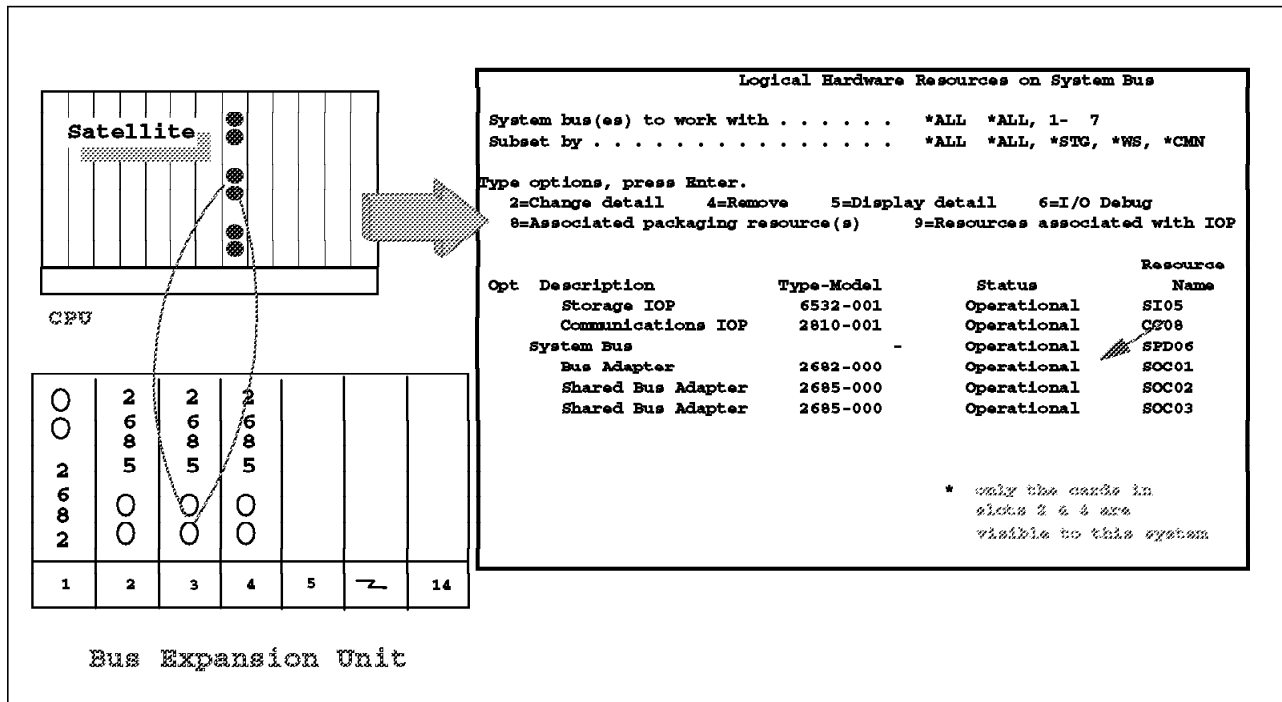


Figure 105. Hardware Service Manager View of OptiConnect Logical Resources on System Bus

Note: A system attached to an OptiConnect card does not show the status of its own card on the bus.

18.5 OptiConnect Satellites

The satellite system has a bus on which it directly addresses all operational cards. In Figure 106, if satellite System A wants to talk to satellite System B, System A's OptiConnect card raises its signal to control the bus. When the bus is available, System A directly sends data to System B by addressing the card. There is no overhead on the hub system since it is not involved in conversations between two satellites.




Figure 106. Satellite to Satellite OptiConnect Connection

18.6 OptiConnect Hub Selection

Selecting which AS/400 system becomes the hub is based primarily on availability requirements. When the hub OptiConnect adapters for all satellites are located on the same bus (a shared bus), the satellites communicate with the hub and other satellites. This assumes that the power for the system expansion unit of the shared bus is switched on. If the hub's system expansion unit loses power, satellite-to-satellite communication is not possible.

When the number of satellites exceeds the maximum number of OptiConnect adapters allowed in a system expansion unit, use a second or third system expansion unit. All satellites can communicate with the hub when their OptiConnect adapters are located on different buses (system expansion units). However, each satellite can only communicate with those satellites that are located in the same system expansion unit (with the same bus number). That means, choose the server that satisfies the communication requirements of all client systems as your hub.

With single path OptiConnect and three or more AS/400 systems, determine which optical link communication path is most critical (CPU1 to CPU3, CPU1 to CPU2, and so forth). Since client-to-server communication is a vital element, the AS/400 system that acts as the server usually functions as the hub, too.

When client-to-server and client-to-client optical link communication are equally important, use a dual path optical link as described in Figure 107 on page 357 in this chapter.

Note: If the hub's power is switched off, the satellites cannot communicate with one another because the OptiConnect adapters are unavailable.

18.7 Dual Path OptiConnect Overview

When the hub's power is turned off, satellite-to-satellite communication is not possible for single path optical link hardware. As a result, the system maintenance or hardware upgrade activity on the hub can cause an outage of satellite-to-satellite optical link communication.

To avoid this, add a second hub to the cluster to provide redundant optical links. With dual OptiConnect hubs, only one hub needs to have power to allow OptiConnect communication between the other OptiConnect systems.

18.7.1 Satellite Dual Path

With dual path OptiConnect, each satellite has two paths (as shown in Figure 107 on page 357)—one path to each hub. Each hub contains one OptiConnect adapter for each satellite. If a hub is unavailable, the satellite still communicates with the remaining hub and every other satellite.




Figure 107. Dual Path OptiConnect Connection

The extra fiber cables and OptiConnect adapters of a dual path connection can appear confusing. When you look at each satellite separately, the connection is easier to understand. Notice the use of top-to-top and bottom-to-bottom cabling and the redundant link.

18.7.2 Hub Dual Path

In Figure 107, we show an example of dual path connection to satellites. The hubs do not share an optical link connection so OptiConnect communication between the hubs is not possible. If hub-to-hub communication is not required, a satellite dual path connection is acceptable.

To enable OptiConnect communication between hubs, add an optical link connection. In Figure 108 on page 358, OptiConnect dedicates *two* buses to OptiConnect, while Hub 2 dedicates only one bus. However, the hubs still do not have dual path OptiConnect. If the OptiConnect adapter in Hub 2 fails, OptiConnect communication between hubs is lost.




Figure 108. Dual Path Hub OptiConnect Connection

For full dual path OptiConnect, add a redundant connection for Hub 2, as shown in Figure 109. When you implement dual path OptiConnect, satellite-to-satellite *and* satellite-to-hub communication are ensured. The communication occurs even when one hub becomes unavailable or an OptiConnect adapter fails. This allows for the highest degree of availability in an OptiConnect cluster.




Figure 109. Full Dual Path OptiConnect Connection

18.8 OptiConnect Diagrams

The following diagram outlines one of the designs explained in this section. Figure 110 shows an OptiConnect four system single bus setup.

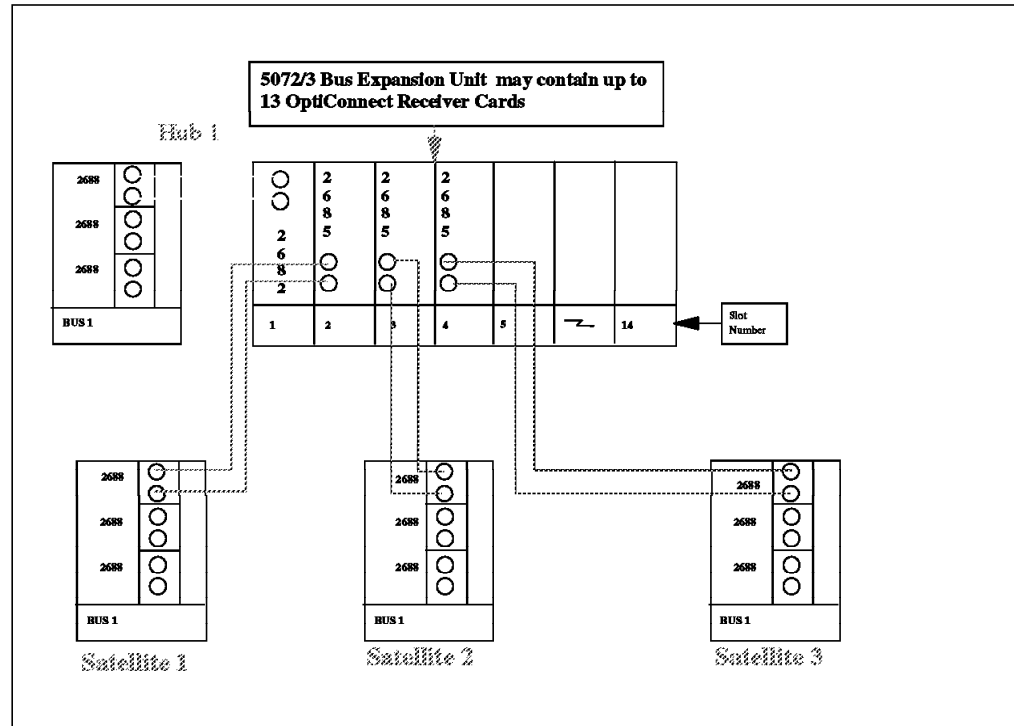


Figure 110. OptiConnect Four System Single Bus

18.9 OptiConnect RPQs

The following chart lists the RPQ numbers required for the OptiConnect for OS/400 solutions described in this chapter.

Table 41 (Page 1 of 2). OptiConnect Feature Codes for 9406 AS/400 Systems

Feature Codes or RPQs	Hardware Description (Bus Speed in Mbps)	Model D, E, F	AS/400 Model 30S, 310, 320	Model 500	Model 510	Model 530, 53S, 6XX	RPQ Number
2669	IMPI Shared Bus Adapter (220)	X	X	X	X	X	843794
2683	RISC Shared Bus Adapter (266)						NA
2685	RISC Shared Bus Adapter (1063)	X	X	X	X	X	843873
2632	5042 Dual Bus Adapter (220)	X	X				843862
2670	5042 Single Bus Adapter (220)	X	X				843805
2680	5044 Single Bus Adapter (266)			X	X	X	843876
2674	Bus Expansion Adapter (266)					X	843896
2673	Bus Expansion Adapter (1063)			X	X		843895
2686	Optical Link Processor (266)			X	X	X	843874
2688	Optical Link Processor (1063)			X	X	X	843875
Bus Units							

Table 41 (Page 2 of 2). OptiConnect Feature Codes for 9406 AS/400 Systems

Feature Codes or RPQs	Hardware Description (Bus Speed in Mbps)	Model D, E, F	AS/400 Model 30S, 310, 320	Model 500	Model 510	Model 530, 53S, 6XX	RPQ Number
5042	IMPI Dual Bus (220- 6+6 slots)	X	X				
5044	Converted 5042 (RISC 266)			X	X	X	
5062	3X0 Single Bus (220 - 13 slots)	RPQ	X				
5070	5X0 Single Bus (266 - 13 slots)			X	X		
5072	530 Single Bus (1063 - 13 slots)					X	
5073	6XX Single Bus (1063 - 13 slots)					X	
Standard HW							
2632	5042 Dual Bus Adapter (220)	X	X				
2684	5044 Dual Bus Adapter			X	X	X	
2670	5062 Single Bus Adapter		X				
2680	5070 Single Bus Adapter (266)			X	X		
2682	5072 Single Bus Adapter (1063)					X	
2673	Bus Expansion Adapter (1063)					X	
2674	Bus Expansion Adapter (266)			X	X		
25XX	Bus Expansion Adapter (220)	X	X				

Note: The feature codes listed are ordered as RPQ numbers for systems prior to V4R2 and some systems after V4R2. The RPQ number matches the feature code listed. After the feature code is available, the equivalent RPQ is withdrawn from marketing.

18.10 For More Information

For details on the installation of OptiConnect products, refer to *OptiConnect for OS/400*, SC41-3414-01 (for V3R2) or SC41-4414-01 (for V3R7).

Appendix A. AS/400 Maximum Capacities

Exceeding system limitations can cause an application or system outage. These limitations can be difficult to predict. However, an administrator can avoid these types of outages by being aware of the system limitations and maximum capacities in advance.

This appendix itemizes some of the capacity limitations and restrictions that can affect the availability of large systems and their applications. For example, an on-line application halts when the size of a file or the number of its members reaches the size limitation. Or, the QSYSOPR message queue becomes full, which can interrupt normal system function and the applications that use this message queue. Refer to Section 10.6, "QSYSOPR Message Queue Wrap When Full" on page 151, for information on how to avoid a full QSYSOPR message queue.

The following tables list the limits or maximum values corresponding to V4R2. Some of these maximum values are different (lower) on prior releases. Also, there are environments or configurations where the actual limit may be less than the stated maximum. For example, certain high-level languages can have more restrictive limits.

Note

The values listed in these tables represent theoretical limits, not thresholds or recommendations. Approaching some of these limits may be unreasonable and can degrade performance. Therefore, practical limits may be lower, depending on system size, configuration, and application environment.

A.1 Limits for Database and SQL

<i>Table 42 (Page 1 of 2). DB2 for AS/400—Database Manager Limits</i>	
DB2 for AS/400—Database Manager Limits	Value
Most columns in a table (number of fields in a record)	8 000
Most columns in a view (number of fields in a record)	8 000
Maximum length of a row including all overhead (number of bytes in a record)	32 766
Maximum size of a table (number of bytes in a database physical file member)	½ TB
Maximum size of an index (number of bytes in an access path)	1 TB
Most rows in a table (number of records in a database physical file member)	2.1 billion
Longest index key (size of key for database files)	2 000
Most columns in an index key (number of key fields in a database file)	120
Most indexes on a table (number of access paths on a database physical file member)	≈ 4 000
Most tables referenced in an SQL statement or a view (number of members that can be joined or number of physical file members in a logical file member)	32
Most host variable declarations in a precompiled program	Amount of storage
Most host variables in an SQL statement ²	≈ 1 024
Longest host variable used for insert or update	32 766
Longest SQL statement	32 767
Most elements in a select list ³	≈ 8 000
Most predicates in a WHERE or HAVING clause	4 690
Maximum number of columns in a GROUP BY clause	120
Maximum total length of columns in a GROUP BY clause	2 000
Maximum number of columns in an ORDER BY clause	10 000
Maximum total length of columns in an ORDER BY clause	10 000
Maximum size of an SQLDA	16 777 215
Maximum number of prepared statements	Amount of storage
Most declared cursors in a program	Amount of storage
Maximum number of cursors opened at one time	Amount of storage
Most tables in a relational database	Amount of storage
Maximum number of constraints on a table	300
Maximum levels allowed for a subselect	32
Maximum length of a comment	2 000

Table 42 (Page 2 of 2). DB2 for AS/400—Database Manager Limits

DB2 for AS/400—Database Manager Limits	Value
Maximum number of rows changed in a unit of work (number of records locked in a single transaction under commitment control)	4 000 000
Maximum number of triggers per physical file	6 triggers
Maximum number of recursive insert and update trigger calls	200
Maximum size of a single journal receiver	2GB
Maximum sequence number for journal entries	2 147 483 136
Maximum number of objects that can be associated with one journal ⁴	250 000
Maximum number of members affected by a single APYJRNCHG command	32 767
Maximum number of remote journal target systems for broadcast mode	255
Maximum number of members in a physical or logical file	32 767
Maximum number of members in a database physical file that can be saved in a single save operation	32 766 (only 16 383 if TYPE(*DATA) and keyed access path)

Table 43 (Page 1 of 2). DB2 for AS/400—SQL Identifier Limits

DB2 for AS/400—SQL Identifier Limits	Value
Longest authorization name	10
Longest correlation name	128
Longest cursor name	18
Longest host identifier	64
Longest server name	18
Longest SQL label	64
Longest statement name	18
Longest unqualified collection name	10
Longest unqualified column name	30
Longest unqualified constraint name	128
Longest unqualified external program name ⁵	10
Longest unqualified nodegroup name	10
Longest unqualified package name	10
Longest unqualified procedure name	128

² The limit is based on the number of host variables that can fit within the longest SQL statement of 32 767 bytes.

³ The limit is based on the size of internal structures generated for the parsed SQL statement.

⁴ This maximum includes physical file members whose changes are currently being journaled, members for which journaling was ended while the current receiver was attached, and journal receivers that are or were associated with the journal while the current journal receiver is attached. If the number of objects is larger than this maximum, journaling does not start.

Table 43 (Page 2 of 2). DB2 for AS/400—SQL Identifier Limits

DB2 for AS/400—SQL Identifier Limits	Value
Longest unqualified specific name	128
Longest unqualified SQL parameter name	64
Longest unqualified SQL variable name	64
Longest unqualified table, view, and index name	128
Unqualified system column name	10
Unqualified system table, view, and index name	10

Table 44. DB2 for AS/400—Numeric Limits

DB2 for AS/400—Numeric Limits	Value
Smallest INTEGER value	-2 147 483 648
Largest INTEGER value	+2 147 483 647
Smallest SMALLINT value	-32 768
Largest SMALLINT value	+32 767
Largest decimal precision	31
Smallest FLOAT value	$\approx -1.79 \times 10^{308}$
Largest FLOAT value	$\approx +1.79 \times 10^{308}$
Smallest positive FLOAT value	$\approx +2.23 \times 10^{-308}$
Largest negative FLOAT value	$\approx -2.23 \times 10^{-308}$
Smallest REAL value	$\approx -3.4 \times 10^{38}$
Largest REAL value	$\approx +3.4 \times 10^{38}$
Smallest positive REAL value	$\approx +1.17 \times 10^{-38}$
Largest negative REAL value	$\approx -1.17 \times 10^{-38}$

Table 45 (Page 1 of 2). DB2 for AS/400—String Limits

DB2 for AS/400—String Limits	Value
Maximum length of CHAR	32 766
Maximum length of VARCHAR	32 740
Maximum length of C NUL-terminated	32 740
Maximum length of GRAPHIC	16 383
Maximum length of VARGRAPHIC	16 370
Maximum length of C NUL-terminated graphic	16 370

⁵ For REXX procedures, the limit is 33.

<i>Table 45 (Page 2 of 2). DB2 for AS/400—String Limits</i>	
DB2 for AS/400—String Limits	Value
Maximum length of character constant	32 740
Maximum length of a graphic constant	16 370
Longest concatenated character string	32 766
Longest concatenated graphic string	16 370

<i>Table 46. DB2 for AS/400—Date and Time Limits</i>	
DB2 for AS/400—Date and Time Limits	Value
Smallest DATE value	0001-01-01
Largest DATE value	9999-12-31
Smallest TIME value	00:00:00
Largest TIME value	24:00:00
Smallest TIMESTAMP value	0001-01-01 - 00.00.00.000000
Largest TIMESTAMP value	9999-12-31 - 24.00.00.000000

A.2 Limits for Communications

<i>Table 47. General Communications Limits</i>	
General Communications Limits	Value
Maximum number of communications configuration objects that can be varied online at IPL	32 767
Maximum number of communications configuration objects that can be in a “varied on” state	64 926
Recommended maximum number of devices allocated to an interactive or communications subsystem	250 to 300
Maximum number of virtual devices that can be specified as automatically configured (QAUTOVRT system value)	32 500 or *NOMAX
Maximum Communications/LAN hardware capabilities	See <i>AS/400e series System Handbook</i> , GA19-5486

<i>Table 48 (Page 1 of 2). SNA Communications Limits</i>	
SNA Communications Limits	Value
Maximum number of SNA controllers per LAN line plus the Network controller	256

<i>Table 48 (Page 2 of 2). SNA Communications Limits</i>	
SNA Communications Limits	Value
Maximum number of SNA CDs across a Frame Relay's NWI lines	256
Maximum number of lines per Frame Relay NWI	256
Maximum number of logical channels per X.25 line	256
Maximum number of controllers on SDLC multidrop lines	254
Maximum number of active sessions per APPC mode	512
Maximum number of modes per APPC device (or APPN location) ⁶	14
Maximum combined number of APPC devices (in any state) and APPN devices (in varied on state)	25 300
Maximum number of APPN intermediate sessions	9 999
Maximum number of devices per controller	254
Maximum size of APPN local location list	476
Maximum size of APPN remote location list	1 898
Maximum size of Asynchronous network address list	294
Maximum size of Asynchronous remote location list	4 995
Maximum size of Retail pass-through list	450
Maximum size of SNA pass-through group	254

<i>Table 49 (Page 1 of 2). TCP/IP Communications Limits</i>	
TCP/IP Communications Limits	Value
Maximum number of interfaces per line	128
Maximum number of interfaces per system	512
Maximum number of routes per system	65 535
Maximum number of ports for TCP	65 535
Maximum number of ports for UDP	65 535
Maximum TCP receive buffer size	8MB
Maximum TCP send buffer size	8MB
Maximum size of a transmission unit on an interface	16 388
Maximum number of Telnet server jobs	132 000
Maximum number of sessions per Telnet server	20
Default maximum number of socket descriptors per job ⁷	200
Maximum number of socket descriptors on the system ⁸	174 761
Maximum size of database files for FTP	1TB

⁶ An APPN location refers to all devices that have the same values for RMTLOCNAME, RMTNETID, and LCLLOCNAME.

<i>Table 49 (Page 2 of 2). TCP/IP Communications Limits</i>	
TCP/IP Communications Limits	Value
Maximum size of IFS files for FTP	2GB
Maximum number of simultaneous SMTP connections	16 client (sending) and 16 server (receiving) connections
Maximum number of MX addresses attempted for SMTP	2

<i>Table 50. OptiConnect Limits</i>	
OptiConnect Limits	Value
Maximum number of systems that can be connected using OptiConnect	32
Maximum number of logical connection paths that can be established between two systems using OptiConnect	2
Maximum distance between systems that are connected using OptiConnect	500 meters (1063 Mbps) or 2 kilometers (266 Mbps)
Maximum number of active jobs that can communicate with any one system using OptiConnect ⁹	16382
Maximum total number of active jobs on one system that can use OptiConnect ⁹	65532

<i>Table 51. Communications Trace Service Tool Limits</i>	
Communications Trace Service Tool Limits	Value
Maximum amount of storage allocated for a single communications trace buffer	64MB
Maximum total amount of storage allocated for all communications trace buffers	128MB
Maximum number of active traces per multiline IOP on pre-V4R1 IOP hardware (limit is removed with new V4R1 IOP hardware)	2

⁷ Default can be changed with DosSetRelMaxFH()—Change the Maximum Number of File Descriptors (see *OS/400 UNIX-Type APIs* in the AS/400 Softcopy Library).

⁸ Actual number may be lower if NFS is used on the system.

⁹ The following count as jobs toward OptiConnect job limits: DDM/DRDA source jobs (user jobs), DDM/DRDA target jobs on server, DB2 multisystem system jobs, APPC controllers using OptiConnect (type *OPC, count as 2 jobs for each controller), jobs using ObjectConnect over OptiConnect, jobs using OptiMover API, and active Remote Journals. Some of these uses are transient for the duration of a function (for example, ObjectConnect SAVRSTxxx) and some are more long term (for example, DDM conversations until reclaimed by RCLDDMCNV or ending the job).

A.3 Limits for Work Management and Security

Table 52. Work Management Limits

Work Management Limits	Value
Maximum number of jobs on the system	163 520
Maximum number of jobs in a subsystem	32 767
Maximum number of prestart jobs initially started when subsystem started	1 000
Maximum number of spooled files per job	9 999
Maximum amount of temporary auxiliary storage that can be specified for a job	1GB or *NOMAX
Maximum number of active memory storage pools	16

Table 53. Security Limits

Security Limits	Value
Maximum number of entries for a user profile ^{10 11 12}	1 048 574
Maximum number of private authorities a user profile can have to successfully save the profile using SAVSYS or SAVSECDTA commands ¹²	≈ 200 000
Maximum number of objects that can be secured by an authorization list	2 097 070
Maximum number of private authorities to an authorization list ¹³	1 048 573
Maximum number of entries in a validation list	2 147 483
Maximum number of user profiles on a system	≈ 340 000

¹⁰ A user profile contains four categories of entries: 1) every object owned by the profile, 2) every private authority the profile has to other objects, 3) every private authority other profiles have to objects owned by this profile, and 4) every object for which this profile is the primary group. The sum of these categories equals the total number of entries for the profile.

¹¹ OS/400 maintains internal user profiles that own objects that are shared or cannot be assigned to a single individual user (for example, QDBSHR owns shared database objects such as database formats, access paths, and so on). These internal user profiles are subject to the same limits as any other user profile on the system.

¹² Using authorization lists or group profiles reduces the number of private authorities and helps avoid this limit (see *Security - Reference*, SC41-5302).

¹³ Limit is due to the maximum number of entries allowed for the user profile that owns the authorization list (one less because a category 01 entry is used for the ownership of the authorization list).

A.4 Limits for Save and Restore

<i>Table 54. Save and Restore Limits</i>	
Save and Restore Limits	Value
Maximum number of related internal objects that can be saved in a single save operation ¹⁴	32 766
Maximum number of members in a database physical file that can be saved in a single save operation	32 766 (only 16 383 if TYPE(*DATA) and keyed access path)
Maximum number of database files in a library that can be saved in a single save operation	≈ 26 200
Maximum number of private authorities a user profile can have to successfully save the profile using SAVSYS or SAVSECDTA commands ¹²	≈ 200 000
Maximum number of names in a save or restore command specifying which objects or libraries to include or exclude in the save or restore operation ¹⁵	300
Maximum number of concurrent save or restore operations	Limited only by available machine resources
Maximum size of an object that can be saved	≈ 1TB
Maximum size of a save file	256GB

¹⁴ Some examples of related objects are:

- All database file objects in a library that are related to each other by dependent logical files.
- All database file objects in a library that are journaled to the same journal when using the save-while-active function.
- All objects in a library when SAVACT(*LIB) is specified.
- All objects in a library when saving to a diskette device.

For most object types, one internal object is saved for each OS/400 object. Some exceptions are:

- Subsystem descriptions:
 - 9 internal objects per subsystem description.
- Database files:
 - At least 1 internal object per physical file member.
 - At least 2 internal objects per member for physical files of TYPE(*DATA) with keyed access paths or constraints.
 - At least 1 internal object per dependent logical file member when ACCPTH(*YES) is specified.

¹⁵ Using generic names to specify groups of objects or libraries can help avoid this limit.

A.5 Miscellaneous Limits

<i>Table 55. Miscellaneous Limits</i>	
Miscellaneous Limits	Value
Maximum number of libraries in a library list	43
Maximum number of libraries in the user portion of the library list	25
Maximum number of objects in a library	≈ 360 000
Maximum number of documents (DLOs) in a user ASP	≈ 349 000
Maximum number of folders in an ASP	65 536
Maximum number of auxiliary storage pools (ASPs)	1 system ASP and 15 user ASPs
Maximum number of objects on the system or in one / root file system	Amount of storage
Maximum number of directories in one directory in the root, QopenSys, or user-defined file systems	32 000 (≈ 100 for good performance)
Maximum size of a stream file	2GB
Maximum size of a user space	16 776 704 bytes
Maximum system and I/O hardware configurations and capacities	See <i>AS/400e series System Handbook</i> , GA19-5486
Maximum number of DASD arms	596
Minimum number of DASD arms required for acceptable performance ¹⁶	Contact your IBM technical representative
Maximum size of QSYSOPR message queue ¹⁷	16MB (≈ 75 000 messages)
Maximum number of input fields that can be specified for a display file	256
Maximum size of files when filing OfficeVision mail locally	16MB

¹⁶ IBM employees should refer to current guidelines contained in *AS4ARMCT PACKAGE* on MKTTOOLS.

¹⁷ When the QSYSOPR message queue gets full, message CPF2460 is issued that states the QSYSOPR message queue could not be extended. PTFs SF47027 (V4R1) and SF45613 (V4R2) are available to avoid this situation and allow the QSYSOPR message queue to wrap. Refer to the PTF cover letters for special instructions. This fix will be incorporated into the next release of the operating system. Contact your software service provider if you have questions.

Appendix B. Evaluating the Time to IPL

IBM has made many improvements to the AS/400 IPL process over the last several releases of OS/400. The average time for a normal IPL for larger customers has been reduced from one or more hours to approximately 15 minutes. Times vary based on many factors, as described in Chapter 4, "IPL Improvements for Availability" on page 33.

As IBM continues to make design changes to lessen the time to IPL, there will be a time when the most benefit comes where the user has an influence (not the design of the IPL) to keep IPL time to a minimum. Most of the tasks the user can control are system maintenance tasks, such as keeping the number of jobs and object descriptions to a minimum.

This appendix introduces a tool (the QWCCRTEC API) that can be used to help you evaluate the time an IPL takes on your system. From the list produced when running the API, you can evaluate what comprises IPL time on your system and determine the steps to improve your IPL based on where the time is spent.

Note that the API used as the tool is an API that is presently shipped on all AS/400 systems. Although used by IBM engineers and some customers to evaluate the IPL process, it is provided on a best effort basis, but is not supported. This means that documentation on how it works, what it produces, and a description of the output is not available. The API may not continue to be provided in future releases of OS/400.

On the command line, enter:

```
CALL QWCCRTEC
```

QWCCRTEC produces a QPSRVDMP spooled file. A sample report follows in Figure 111 on page 372. The SRC codes that mark the stage of the IPL process are highlighted with the date and time stamp noted in the right hand column. Refer to *AS/400 Licensed Internal Code Diagnostic Aids — Volume 1*, LY44-5900, for a further description of the SRCs and the IPL process.

OBJECT TYPE-	SPACE								*QTSP			
NAME-	QWCSRCDATAOUTPUT								19	SUBTYPE-	EF	
CREATION-	04/08/98 06:06:45								SIZE-	0000004000		
ATTRIBUTES-	0000								ADDRESS-	D812DBA185 000000		
SPACE ATTRIBUTES-												
000000	00FFFF00	00000060	19EFD8E6	C3E2D9C3	C4C1E3C1	D6E4E3D7	E4E34040	40404040	*	-	QWCSRCDATAOUTPUT	*
000020	40404040	40404040	40020000	00000000	00003F00	00010000	00000000	00000000	*			*
000040	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	*			*
000000	E7D7C640	D7E6D9C4	E6D54040	404040F0	F361F2F7	61F9F840	F0F47AF0	F07AF0F2	*	XPF PWRDWN	03/27/98 04:00:02	
000020	0000C4F9	F0F000F2	F7F4F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F07AF0F2	*	D900 2740	03/27/98 04:00:02*	
000040	0000C4F9	F0F000F2	F7F5F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F17AF5F7	*	D900 2750	03/27/98 04:01:57*	
000060	0000C4F9	F0F000F2	F7F7F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F17AF5F7	*	D900 2770	03/27/98 04:01:57*	
000080	0000C4F9	F0F000F2	F7F8F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F17AF5F8	*	D900 2780	03/27/98 04:01:58*	
0000A0	0000C4F9	F0F000F2	F8F1F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F17AF5F8	*	D900 2790	03/27/98 04:01:58*	
0000C0	0000C4F9	F0F000F2	F7C3F000	000000F0	F361F2F7	61F9F840	F0F47AF0	F17AF5F8	*	D900 27C0	03/27/98 04:01:58*	
0000E0	C5958440	D7E6D9C4	E6D54040	404040F0	F361F2F7	61F9F840	F0F47AF0	F27AF0F2	*	End PWRDWN	03/27/98 04:02:02	
000100	E7D7C640	C9D7D340	40404040	404040F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F6	*	XPF IPL	03/27/98 04:10:26	
000120	0000C3F9	F0F000F2	F8F1F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F6	*	C900 2810	03/27/98 04:10:26	
000140	0000C3F9	F0F000F2	F8F2F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F6	*	C900 2820	03/27/98 04:10:26	
000160	00000000	F1F000F2	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 20 0010	03/27/98 04:10:27*	
000180	00000000	F1F000F2	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 20 0020	03/27/98 04:10:27*	
0001A0	00000000	F1F000F3	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 30 0010	03/27/98 04:10:27*	
0001C0	00000000	F1F000F3	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 30 0020	03/27/98 04:10:27*	
0001E0	00000000	F1F000F3	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 30 0030	03/27/98 04:10:27*	
000200	00000000	F1F000F3	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 30 0040	03/27/98 04:10:27*	
000220	00000000	F1F000F3	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 30 0050	03/27/98 04:10:27*	
000240	00000000	F1F000F3	C100F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0010	03/27/98 04:10:27*	
000260	0000C3F9	F0F000F2	F8F3F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	C900 2830	03/27/98 04:10:27	
000280	00000000	F1F000F3	C100F0F0	F1F800F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0018	03/27/98 04:10:27*	
0002A0	00000000	F1F000F3	C100F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0020	03/27/98 04:10:27*	
0002C0	00000000	F1F000F3	C100F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0030	03/27/98 04:10:27*	
0002E0	00000000	F1F000F3	C100F0F0	F3F500F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0035	03/27/98 04:10:27*	
000300	00000000	F1F000F3	C100F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0050	03/27/98 04:10:27*	
000320	00000000	F1F000F3	C100F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0060	03/27/98 04:10:27*	
000340	00000000	F1F000F3	C100F0F0	F7F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0070	03/27/98 04:10:27*	
000360	00000000	F1F000F3	C100F0F0	F8F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0080	03/27/98 04:10:27*	
000380	00000000	F1F000F3	C100F0F0	F8F200F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0082	03/27/98 04:10:27*	
0003A0	00000000	F1F000F3	C100F0F0	F9F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0090	03/27/98 04:10:27*	
0003C0	00000000	F1F000F3	C100F0F0	F9F800F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 0098	03/27/98 04:10:27*	
0003E0	00000000	F1F000F3	C100F0F0	C1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00A0	03/27/98 04:10:27*	
000400	00000000	F1F000F3	C100F0F0	C1F400F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00A4	03/27/98 04:10:27*	
000420	00000000	F1F000F3	C100F0F0	C1F800F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00A8	03/27/98 04:10:27*	
000440	00000000	F1F000F3	C100F0F0	C1C100F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00AA	03/27/98 04:10:27*	
000460	00000000	F1F000F3	C100F0F0	C1C300F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00AC	03/27/98 04:10:27*	
000480	00000000	F1F000F3	C100F0F0	C2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00B0	03/27/98 04:10:27*	
0004A0	00000000	F1F000F3	C100F0F0	C2F400F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00B4	03/27/98 04:10:27*	
0004C0	00000000	F1F000F3	C100F0F0	C2F800F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00B8	03/27/98 04:10:27*	
0004E0	00000000	F1F000F3	C100F0F0	C2C100F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F7	*	10 3A 00BA	03/27/98 04:10:27*	
000500	00000000	F1F000F3	C100F0F0	C3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3A 00C0	03/27/98 04:10:28*	
000520	00000000	F1F000F3	C100F0F0	C4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3A 00D0	03/27/98 04:10:28*	
000540	00000000	F1F000F3	C100F0F0	C4F400F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3A 00D4	03/27/98 04:10:28*	
000560	00000000	F1F000F3	C100F0F0	C5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3A 00E0	03/27/98 04:10:28*	
000580	00000000	F1F000F3	C100F0F0	C6F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3A 00F0	03/27/98 04:10:28*	
0005A0	0000C3F9	F0F000F2	F8C3F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	C900 28C0	03/27/98 04:10:28	

Figure 111 (Part 1 of 7). Sample Report

0005C0	00000000	F1F000F3	C600F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 3F 0010	03/27/98	04:10:28*
0005E0	0000C3F9	F0F000F2	F8F2F500	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	C900 2825	03/27/98	04:10:28
000600	0000C3F9	F0F000F2	F8C3F500	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	C900 28C5	03/27/98	04:10:28
000620	00000000	F1F000F4	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 40 0010	03/27/98	04:10:28*
000640	00000000	F1F000F4	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 40 0020	03/27/98	04:10:28*
000660	00000000	F1F000F4	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 40 0050	03/27/98	04:10:28*
000680	00000000	F1F000F5	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0010	03/27/98	04:10:28*
0006A0	00000000	F1F000F5	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0020	03/27/98	04:10:28*
0006C0	00000000	F1F000F5	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0030	03/27/98	04:10:28*
0006E0	00000000	F1F000F5	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0040	03/27/98	04:10:28*
000700	00000000	F1F000F5	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0050	03/27/98	04:10:28*
000720	00000000	F1F000F5	F000F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 50 0060	03/27/98	04:10:28*
000740	00000000	F1F000F6	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 60 0010	03/27/98	04:10:28*
000760	00000000	F1F000F6	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 60 0020	03/27/98	04:10:28*
000780	00000000	F1F000F7	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 70 0010	03/27/98	04:10:28*
0007A0	00000000	F1F000F7	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 80 0010	03/27/98	04:10:28*
0007C0	00000000	F1F000F8	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 80 0030	03/27/98	04:10:28*
0007E0	00000000	F1F000F8	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 80 0040	03/27/98	04:10:28*
000800	00000000	F1F000F8	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 80 0050	03/27/98	04:10:28*
000820	00000000	F1F000F8	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 A0 0010	03/27/98	04:10:28*
000840	00000000	F1F000C1	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 A0 0020	03/27/98	04:10:28*
000860	00000000	F1F000C2	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F8	*	10 B0 0010	03/27/98	04:10:28*
000880	00000000	F1F000C2	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F9	*	10 B0 0040	03/27/98	04:10:29*
0008A0	00000000	F1F000C2	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F9	*	10 B0 0050	03/27/98	04:10:29*
0008C0	00000000	F1F000C4	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF2F9	*	10 D0 0010	03/27/98	04:10:29*
0008E0	00000000	F1F000C6	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0010	03/27/98	04:10:30*
000900	00000000	F1F000C6	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0020	03/27/98	04:10:30*
000920	00000000	F1F000C6	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0030	03/27/98	04:10:30*
000940	00000000	F1F000C6	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0040	03/27/98	04:10:30*
000960	00000000	F1F000C6	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0050	03/27/98	04:10:30*
000980	00000000	F1F000C6	F000F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	10 F0 0060	03/27/98	04:10:30*
0009A0	00000000	F2F000F1	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	20 10 0010	03/27/98	04:10:30*
0009C0	00000000	F2F000F1	F000F0F0	F1F500F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F0	*	20 10 0015	03/27/98	04:10:30*
0009E0	00000000	F2F000F1	F000F0F0	F1F800F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0018	03/27/98	04:10:35*
000A00	00000000	F2F000F1	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0020	03/27/98	04:10:35*
000A20	00000000	F2F000F1	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0030	03/27/98	04:10:35*
000A40	00000000	F2F000F1	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0040	03/27/98	04:10:35*
000A60	00000000	F2F000F1	F000F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0050	03/27/98	04:10:35*
000A80	00000000	F2F000F1	F000F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0060	03/27/98	04:10:35*
000AA0	00000000	F2F000F1	F000F0F0	F7F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 10 0070	03/27/98	04:10:35*
000AC0	00000000	F2F000F1	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	20 18 0010	03/27/98	04:10:35*
000AE0	0000C3F9	F0F000F2	F9F1F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F5	*	C900 2910	03/27/98	04:10:35
000B00	00000000	F2F000F1	F800F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F7	*	20 18 0030	03/27/98	04:10:37*
000B20	00000000	F2F000F1	F800F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F7	*	20 18 0040	03/27/98	04:10:37*
000B40	00000000	F2F000F2	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F7	*	20 20 0010	03/27/98	04:10:37*
000B60	00000000	F2F000F2	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F7	*	20 28 0010	03/27/98	04:10:37*
000B80	0000C3F9	F0F000F2	F9F2F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F7	*	C900 2920	03/27/98	04:10:37
000BA0	00000000	F2F000F2	F900F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F8	*	20 29 0010	03/27/98	04:10:38*
000BC0	0000C3F9	F0F000F2	F9F2F500	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF3F8	*	C900 2925	03/27/98	04:10:38
000BE0	00000000	F2F000F3	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F5	*	20 30 0010	03/27/98	04:10:45*
000C00	0000C3F9	F0F000F2	F9F3F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F5	*	C900 2930	03/27/98	04:10:45
000C20	00000000	F2F000F3	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F6	*	20 30 0020	03/27/98	04:10:46*
000C40	00000000	F2F000F3	F400F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F7	*	20 34 0010	03/27/98	04:10:47*
000C60	00000000	F2F000F3	F400F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F7	*	20 34 0020	03/27/98	04:10:47*
000C80	00000000	F2F000F5	F100F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F8	*	20 51 0010	03/27/98	04:10:48*
000CA0	0000C3F9	F0F000F2	F9F6F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F8	*	C900 2960	03/27/98	04:10:48
000CC0	00000000	F2F000F5	F100F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F8	*	20 51 0020	03/27/98	04:10:48*
000CE0	00000000	F2F000F5	F100F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F8	*	20 51 0040	03/27/98	04:10:48*

Figure 111 (Part 2 of 7). Sample Report

000D00	0000C3F9	F0F000F2	F9F6F500	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF4F8	*	C900 2965	03/27/98	04:10:48
000D20	00000000	F2F000F5	F100F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F07AF5F5	*	20 51 0050	03/27/98	04:10:55*
000D40	0000C3F9	F0F000F2	C1F8F500	000000F0	F361F2F7	61F9F840	F0F47AF1	F07AF5F5	*	C900 2A85	03/27/98	04:10:55
000D60	0000C3F9	F0F000F2	F9F6F700	000000F0	F361F2F7	61F9F840	F0F47AF1	F17AF2F1	*	C900 2967	03/27/98	04:11:21
000D80	00000000	F2F000F5	F100F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F17AF2F1	*	20 51 0060	03/27/98	04:11:21*
000DA0	00000000	F2F000F5	F100F0F0	F8F000F0	F361F2F7	61F9F840	F0F47AF1	F37AF5F8	*	20 51 0080	03/27/98	04:13:58*
000DC0	0000C3F9	F0F000F2	C1F8F700	000000F0	F361F2F7	61F9F840	F0F47AF1	F37AF5F8	*	C900 2A87	03/27/98	04:13:58
000DE0	00000000	F2F000F5	F100F0F0	F9F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 51 0090	03/27/98	04:14:01*
000E00	0000C3F9	F0F000F2	F9F6F800	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	C900 2968	03/27/98	04:14:01
000E20	00000000	F2F000F5	F100F0F0	C1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 51 00A0	03/27/98	04:14:01*
000E40	00000000	F2F000F5	F100F0F0	C2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 51 00B0	03/27/98	04:14:01*
000E60	00000000	F2F000F5	F100F0F0	C3F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 51 00C0	03/27/98	04:14:01*
000E80	00000000	F2F000F5	F700F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 57 0010	03/27/98	04:14:01*
000EA0	00000000	F2F000F5	F700F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 57 0020	03/27/98	04:14:01*
000EC0	00000000	F2F000F5	F700F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 57 0040	03/27/98	04:14:01*
000EE0	00000000	F2F000F5	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 60 0010	03/27/98	04:14:01*
000F00	00000000	F2F000F6	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 60 0030	03/27/98	04:14:01*
000F20	00000000	F2F000F6	F000F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 60 0040	03/27/98	04:14:01*
000F40	00000000	F2F000F6	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 68 0010	03/27/98	04:14:01*
000F60	00000000	F2F000F7	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 70 0010	03/27/98	04:14:01*
000F80	00000000	F2F000F7	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 78 0010	03/27/98	04:14:01*
000FA0	00000000	F2F000F7	F800F0F0	F1F800F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 78 0018	03/27/98	04:14:01*
000FC0	00000000	F2F000F7	F800F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 78 0020	03/27/98	04:14:01*
000FE0	00000000	F2F000F7	F800F0F0	F2F200F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F1	*	20 78 0022	03/27/98	04:14:01*
001000	00000000	F2F000F8	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F2	*	20 80 0010	03/27/98	04:14:02*
001020	0000C3F9	F0F000F2	F9F7F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF0F2	*	C900 2970	03/27/98	04:14:02
001040	00000000	F2F000F8	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 88 0010	03/27/98	04:14:48*
001060	0000C3F9	F0F000F2	F9F8F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	C900 2980	03/27/98	04:14:48
001080	00000000	F2F000F8	F800F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 88 0020	03/27/98	04:14:48*
0010A0	00000000	F2F000F8	F800F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 88 0040	03/27/98	04:14:48*
0010C0	00000000	F2F000F8	F800F0F0	F5F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 88 0050	03/27/98	04:14:48*
0010E0	00000000	F2F000F9	F800F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 98 0020	03/27/98	04:14:48*
001100	0000C3F9	F0F000F2	F9C1F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	C900 29A0	03/27/98	04:14:48
001120	00000000	F2F000F9	F800F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 98 0030	03/27/98	04:14:48*
001140	00000000	F2F000F9	F800F0F0	F4F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 98 0040	03/27/98	04:14:48*
001160	00000000	F2F000F9	F800F0F0	F6F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 98 0060	03/27/98	04:14:48*
001180	00000000	F2F000F9	F800F0F0	F7F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 98 0070	03/27/98	04:14:48*
0011A0	00000000	F2F000C1	F000F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	20 A0 0010	03/27/98	04:14:48*
0011C0	0000C3F9	F0F000F2	F9C2F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF4F8	*	C900 29B0	03/27/98	04:14:48
0011E0	00000000	F2F000C1	F000F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF5F5	*	20 A0 0020	03/27/98	04:14:55*
001200	00000000	F2F000C1	F000F0F0	F3F000F0	F361F2F7	61F9F840	F0F47AF1	F47AF5F9	*	20 A0 0030	03/27/98	04:14:59*
001220	0000C3F9	F0F000F2	F9C3F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F47AF5F9	*	C900 29C0	03/27/98	04:1 4:5
001240	00000000	F2F000C1	F000F0F0	F3F100F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F5	*	20 A0 0031	03/27/98	04:15:45*
001260	0000C3F9	F0F000F2	C1F8F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F5	*	C900 2A80	03/27/98	04:15:45
001280	00000000	F2F000C1	F100F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F5	*	20 A1 0010	03/27/98	04:15:45*
0012A0	00000000	F2F000C1	F100F0F0	F2F000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F5	*	20 A1 0020	03/27/98	04:15:45*
0012C0	00000000	F2F000C1	F400F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F6	*	20 A4 0010	03/27/98	04:15:46*
0012E0	0000C3F9	F0F000F2	C1F9F000	000000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F6	*	C900 2A90	03/27/98	04:15:46
001300	00000000	F2F000C1	F800F0F0	F0F500F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F6	*	20 A8 0005	03/27/98	04:15:46*
001320	00000000	F2F000C1	F800F0F0	F0F600F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F6	*	20 A8 0006	03/27/98	04:15:46*
001340	00000000	F2F000C1	F800F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F6	*	20 A8 0010	03/27/98	04:15:46*
001360	00000000	F2F000C1	F800F0F0	F1F500F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F8	*	20 A8 0015	03/27/98	04:15:48*
001380	00000000	F2F000C1	F900F0F0	F1F500F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F8	*	20 A9 0015	03/27/98	04:15:48*
0013A0	00000000	F2F000C1	F900F0F0	F1F000F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F8	*	20 A9 0010	03/27/98	04:15:48*
0013C0	00000000	F2F000C1	F900F0F0	F3F100F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F8	*	20 A9 0031	03/27/98	04:15:48*
LINES 0013E0 TO 00143F SAME AS ABOVE												
001440	00000000	F2F000C1	F900F0F0	F3F100F0	F361F2F7	61F9F840	F0F47AF1	F57AF4F9	*	20 A9 0031	03/27/98	04:15:49*
LINES 001460 TO 00149F SAME AS ABOVE												

Figure 111 (Part 3 of 7). Sample Report


```

002A60 00000000 F2F000C3 F000F0F0 F5F800F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F3 * 20 C0 0058 03/27/98 04:16:33*
002A80 00000000 F2F000C3 F000F0F0 F6F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F3 * 20 C0 0060 03/27/98 04:16:33*
002AA0 00000000 F2F000C3 F000F0F0 F9F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 C0 0090 03/27/98 04:16:34*
002AC0 00000000 F2F000C4 F000F0F1 F0F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D0 0100 03/27/98 04:16:34*
002AE0 0000C3F9 F0F000F2 C3F4F000 000000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * C900 2C40 03/27/98 04:16:34
002B00 00000000 F2F000C4 F000F0F8 F1F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D0 0810 03/27/98 04:16:34*
002B20 00000000 F2F000C4 F800F0F0 F2F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D8 0020 03/27/98 04:16:34*
002B40 0000C3F9 F0F000F2 C3F2F000 000000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * C900 2C20 03/27/98 04:16:34
002B60 00000000 F2F000C4 F800F0F0 F2F300F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D8 0023 03/27/98 04:16:34*
002B80 0000C3F9 F0F000F2 C3F2F500 000000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * C900 2C25 03/27/98 04:16:34
002BA0 00000000 F2F000C4 F800F0F0 F2F500F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D8 0025 03/27/98 04:16:34*
002BC0 00000000 F2F000C4 F800F0F0 F2F700F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F4 * 20 D8 0027 03/27/98 04:16:34*
002BE0 0000C3F9 F0F000F2 C6F0F000 000000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F7 * C900 2F00 03/27/98 04:16:37
002C00 00000000 F2F000C4 F800F0F8 F2F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F7 * 20 D8 0820 03/27/98 04:16:37*
002C20 C5958440 968640C9 D7D34040 404040F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F7 * End of IPL 03/27/98 04:16:37
002C40 00000000 F2F000C4 F900F0F0 F1F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F7 * 20 D9 0010 03/27/98 04:16:37*
002C60 00000000 F2F000C4 F900F0F0 F2F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF3F7 * 20 D9 0020 03/27/98 04:16:37*
002C80 00000000 F2F000C4 F900F0F0 F3F000F0 F361F2F7 61F9F840 F0F47AF1 F67AF4F6 * 20 D9 0030 03/27/98 04:16:46*

```

Figure 111 (Part 7 of 7). Sample Report

If you run this tool on a regular basis, you can use the output to estimate how long the IPL remains at each phase.

Recommendation

Run the QWCCRTEC tool periodically to keep track of both normal and abnormal IPL times. An analysis of where time is spent can reveal where you can make improvements in the time to IPL your system.

For more information, the SRCs associated with each IPL phase are listed in Section 4.6, "Marking the Progress of an IPL with SRC Codes" on page 44.

Appendix C. Save and Restore Rates of IBM Tape Drives for Sample Workloads

As you plan for offline media to handle your save and restore requirements, use the information in this appendix as a reference. Note that the throughput that we obtained was measured for the workload described in dedicated mode. Results on your system may not be the same as what is represented here. However, the figures should be relative and useful as you plan to minimize the time for saving and restoring your system.

The following workloads were designed to help evaluate the performance of save and restore operations at the same time. Familiarization with the makeup of the workloads helps clarify the differences in the save and restore rates presented in Chapter 8, "Save, Restore, and System Performance for Availability" on page 95. It also helps you understand how they may apply to your environment.

USERENV The user environment workload (USERENV) consists of four libraries:

- The first library contains four source files (for a total of 1 204 members) that comprise about 39MB of space.
- The second library consists of 28 database files, ranging in size from 2MB to 200MB, for a total of 470MB in size.
- The third library consists of 200 program objects, with an average size of about 100KB, for a total size of 20MB.
- The fourth library is 12MB in size and consists of 2 156 objects of various types.

Therefore, the USERENV workload consists of about 556MB.

SRC The SRC workload consists of:

Four source files that are in the first library of the USERENV. These source files occupy about 39MB of space and contain a total of 1,204 members.

200MB The 200MB workload is:

A single member database file that is about 200MB in size. This workload is saved using the Save Object (SAVOBJ) command and restored using the Restore Object (RSTOBJ) command.

2GB The 2GB workload is:

A single member database file that is about 2GB in size. This workload is saved using the SAVOBJ command and restored using the RSTOBJ command.

DLO The DLO workload consists of:

Eight folders with 3,700 documents in them. The documents range in size from 53KB to 233KB for a combined size of 396MB. All of the documents reside in the first level folder structure.

integrated file system

The type of data stored in the integrated file system has been changing since it was made available in OS/400 V3R1. In the past, the type of data stored in the file system mainly consisted of client programs. Programs do not compact or compress so they are saved

or restored at the basic rate of the tape drive that is used. With the introduction of Lotus Notes and Web functions, more files containing data are stored in the integrated 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.

On a system with an even mixture of client programs, for example, Lotus Notes databases and Web home pages should save and restore in the range of the USERENV workload described in our charts. If the data stored on the system is largely made up of database files, the save and restore rates probably reflect the 200MB or larger file-workload type, depending on the size and number of database files. If the data is largely made up of Web files, the save rates reflect those of the USERENV workload down to the SRC workload. You may see similar save and restore rates, as demonstrated in this example, depending upon your data and its compaction capabilities. (Web files tend to be numerous but contain many small HTML files such as small home pages.)

Keep in mind that 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 composition reverses the previous description.

In all situations, the actual data dictates the save and restore rates. Be familiar with the type of data on your system to estimate save and restore rates for your environment.

The **concurrent save and restore** workload offers the ability to save and restore a single library to multiple tape drives at the same time, a function that is available at V4R2. You can generally expect to save 1.8 times as much data than when using one tape drive in the same amount of time.

The workloads used for this test were two 2GB files in the same library and the USERENV workload. All objects from the four USERENV libraries were combined into one library, and all of the objects were duplicated. The duplicate objects were named so that the original objects were saved to one tape drive and the duplicate objects were saved to another.

Note: Most of the save and restore rates were obtained from a restricted state when all subsystems were brought to an end status using the ENDSBS(*ALL) command. The workloads for concurrent saves were on a dedicated system. On a dedicated system, as opposed to a restricted system, the QBATCH subsystem runs so concurrent jobs can process.

C.1 Comparing Performance Data

When comparing the performance data in this appendix (and Chapter 8, “Save, Restore, and System Performance for Availability” on page 95) with the actual performance on your system, remember that the performance of save and restore operations *is data dependent*. If the same save or restore operation is performed on data from three different systems, three different rates result, even using similar tape drives. The performance fluctuation is inherent to differences in the data itself. The same applies to the configuration of the DASD units on which the data is stored. The more DASD units installed on the system, the better the performance.

Generally speaking, the large file data used in this test scenario compacted around a 3:1 ratio. The following formula illustrates how performance ratings are obtained:

$$((\text{TapeSpeed} * \text{LossFromWorkLoadType}) * \text{Compaction}) = \text{MB/Sec}$$

We then multiply MB/Sec * 3600 = MB/HR

The reality of this formula is that the LossFromWorkLoadType is far more complex than described here. The different workloads have different overheads and compaction rates. Plus, the drives use different buffer sizes and compaction algorithms. The attempt is to group these workloads as examples of what might happen with a certain type of drive and workload.

Note: The previous formulas and following charts are designed to give you an idea about what save and restore rates you can achieve from a particular IBM tape drive. Your results may differ. Because your data is unique to your company, the best tape solution for *your requirements must take into account many different factors from the USERENV and USR environment* prescribed in the IBM tests described in this appendix. Consider these factors when sizing tape requirements for your environment:

1. System size
2. Model of tape drive
3. Number of tapes that are required
4. Whether you are performing an attended or unattended save operation

C.2 Lower Speed Tape Drives

Lower performing tape drives (for example, the 6380 and 6390) are a limiting factor in save and restore throughput. The save rates are approximately the same, regardless of the size system they are attached to. For testing purposes on the 6380, compression was substituted for compaction. To determine tape speed, we used the formula:

$$\text{TapeSpeed} * \text{LossFromWorkLoadType}$$

Use the following LossFromWorkLoadType approximations (for save operations) in the formula when looking at lower performing tape drives.

Workload Type	Amount of loss
2GB	95%
USERENV, 200M, DLO	85%
SRC	50%

Compression of the 6380 = 1.8, and compaction for the 6390 = 1.8.

$ \begin{array}{ccccccc} \text{TapeSpeed} & * & \text{LossFromWorkLoadType} & * & \text{Compaction} & & \\ 6390 & & .5 & * & .95=(.475) & * & 1.8 & =(.713) \text{ MB/S} * 3600 = 3078 \text{ MB/HR} \end{array} $

Figure 112. 6390 Saving Large Files — Example

C.3 Medium Speed Tape Drives

Overhead is different for medium performing tape drives (for example, the 6385, 3490, and 3570). These tape drives are designed with different technologies, which therefore, makes it difficult to compare them. The 3570 uses optimum block size and a different compaction method than other tape devices. Using USEOPTBLK(*YES) can make the 3570 an efficient drive for systems that are CPU-constrained.

The 3570 also performs better with data compaction. However, if the data cannot be compacted, the 3490 is faster than the 3570. The 3490 also outperforms the 3570 if a large amount of integrated file system data is used. In this case, the 3490 is about 30% faster than the 3570. For general data that can be compacted, the 3570 is better at compaction and outperforms the 3490 by about 10%.

<i>Table 57. LossFromWorkLoadType Approximations (Save Operations)</i>	
Workload Type	Amount of loss
2GB	85%
USERENV, 200M, DLO	65%
SRC	25%

Note: The 3490 does *not* make full use of compaction. The 1.8 factor was derived for simulation purposes only.

	TapeSpeed	*	LossFromWorkLoadType	*	Compaction	
6385	1.5	*	.85=(1.28)	*	1.8	= (2.3) MB/S * 3600 = 8280 MB/HR
3570	2.2	*	.85=(1.87)	*	2.5	= (4.68) MB/S * 3600 = 16848 MB/HR
3490	3.0	*	.85=(2.60)	*	1.8	= (4.42) MB/S * 3600 = 15912 MB/HR

Figure 113. 6385, 3570, and 3490 Saving Large Files — Example

C.4 Highest Speed Tape Drives

The highest performing tape drives are capable of data rates of 3MB or more.

The overhead for the fastest AS/400 tape drive (the 3590) is different from other tape drive types. The 3590 takes advantage of optimum block. Its nine Mbps rating allows large system users to keep their save windows at a minimum.

The 3590 performs best on large files as shown in Table 58 through Table 66 on page 385. See Section 8.1.4, "Save and Restore Tips and Techniques for Better Performance" on page 98 for details on getting the most out of your 3590.

<i>Table 58. LossFromWorkLoadType Approximations (Save Operations)</i>	
Workload Type	Amount of loss
2 GB	95%
USERENV, 200M, DLO	44%
SRC	12%

Note: The 3590 does **not** make full use of compaction. The 1.6 factor was derived for simulation purposes only.

$\text{TapeSpeed} * \text{LossFromWorkLoadType} * \text{Compaction}$ $3590 \quad 9.0 \quad * \quad .95=(8.6) \quad * \quad 1.6 \quad =(13.76) \text{ MB/S} * 3600 = 49536 \text{ MB/HR}$
--

Figure 114. 3590 Saving Large Files — Example

C.5 Save and Restore Rates

The save and restore rates that are quoted in this section are expressed in terms of megabytes per hour (MB/HR). All of the measurements use the options for compression (DTACPR) and compaction (COMPACT) as described in Chapter 8, “Save, Restore, and System Performance for Availability” on page 95. We use the USEOPTBLK parameter where available. It is noted by the “1” after the tape device. All systems measurements are from OS/400 V4R2 systems.

<i>Table 59. Save Rates for Model 400 Feature 2133</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
6380	1310	1740	1800	1890	1780
6390	1350	2500	2100	2600	3190
7208-342 ¹	2340	10200	9500	15500	14990
3570 ¹	4550	11550	14000	20000	14000
¹ The save operation was performed with USEOPTBLK(*YES).					

<i>Table 60. Restore Rates for Model 400 Feature 2133</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
6380	610	1740	1800	1890	1780
6390	850	2000	2500	2600	3190
7208-342 ¹	1610	6840	6560	15640	3520
3570 ¹	2300	8000	17000	19000	3800
¹ The save operation was performed with USEOPTBLK(*YES).					

<i>Table 61 (Page 1 of 2). Save Rates for Model 510 Feature 2144</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
6380	1410	1760	1770	1900	1780
6382	1830	3560	3610	4050	3640
6390	1280	2470	2050	2625	3310
7208-342 ¹	2240	9970	9620	15580	14790

<i>Table 61 (Page 2 of 2). Save Rates for Model 510 Feature 2144</i>					
Tape Device	Workload				
6385	1320	5300	6590	9640	1980
3490	4400	11000	11000	14800	13850
3570 ¹	4500	11800	16000	20600	12400
3590 ¹	6100	18600	25690	44750	24300
1 The save operation was performed with USEOPTBLK(*YES).					

<i>Table 62. Restore Rates for Model 510 Feature 2144</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
6380	1120	1680	1760	1894	1730
6382	1350	2010	2260	2180	2690
6390	1560	2400	2400	2700	3170
7208-342 ¹	1980	8220	12470	15850	4540
6385	1470	2390	7922	9530	2260
3490	2400	8000	11200	14300	4650
3570 ¹	2800	9000	15700	20500	4700
3590 ¹	3100	12100	27700	39900	5010
1 The save operation was performed with USEOPTBLK(*YES).					

<i>Table 63. Save Rates for Model S20 Feature 2166</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
3490	4400	12800	13500	19000	15000
3570 ¹	5800	14100	14900	20800	15000
3590 ¹	6200	24000	30000	49800	25000
1 The save operation was performed with USEOPTBLK(*YES).					

<i>Table 64. Restore Rates for Model S20 Feature 2166</i>					
Tape Device	Workload				
	SRC	USERENV	200M	2GB	DLO
3490	4100	11400	13300	17300	10400
3570 ¹	4500	14400	17000	21000	10500
3590 ¹	4690	23900	37000	49900	20000
1 The save operation was performed with USEOPTBLK(*YES).					

<i>Table 65. Model S20 Feature 2166 Concurrent Saves</i>		
Tape Device	Workload	
	USERENV	2GB
3570 ¹ 1st Drive	12600	19800
3570 ¹ 2nd Drive	13600	20400
3590 ¹ 1st Drive	23500	48800
3590 ¹ 2nd Drive	23200	48900
1 The save operation was performed with USEOPTBLK(*YES).		

<i>Table 66. Model S20 Feature 2166 Concurrent Restores</i>		
Tape Device	Workload	
	USERENV	2GB
3570 ¹ 1st Drive	12500	19400
3570 ¹ 2nd Drive	13200	20000
3590 ¹ 1st Drive	23300	50100
3590 ¹ 2nd Drive	23200	48800
1 The save operation was performed with USEOPTBLK(*YES).		

C.6 Save and Restore Rates for Optical Device

The following save and restore performance measurements were made on an AS/400 Model 510 Feature 2143 using the 3995 Model C48 Optical Library. The save and restore rates are expressed in terms of megabytes per hour (MB/HR). They include the processing time that is required to complete the operation, but not the time that is required for the autochanger to load or unload the optical cartridge.

<i>Table 67. Save and Restore Rates for Optical Device</i>				
Operation	Workload			
	SRC	USERENV	200M	2GB
Save	1340	1925	2060	2160
Restore	830	3090	5910	6470

Appendix D. OptiConnect for OS/400 Terminology and Hardware Overview

When describing an OptiConnect solution, the terminology used can seem unfamiliar to even a seasoned AS/400 administrator. This appendix introduces you to OptiConnect terminology and its hardware components. As in Chapter 18, "OptiConnect for OS/400" on page 351, this section also refers only to RISC models.

D.1 OptiConnect for OS/400 Terminology

To understand OptiConnect for OS/400, you must first understand the basic terminology. This section explains such fundamental terms as cluster, satellite, hub, link, path, and redundancy.

D.1.1 An OptiConnect Cluster

An OptiConnect cluster is a collection of AS/400 systems connected by dedicated fiber optic system bus cables. The AS/400 Model 500, 510, and 50S have one internal system bus and up to six external optical system buses. The AS/400 Models 530 and 53S have one internal system bus and up to 18 external optical system buses. Each *external* optical bus is available for use by OptiConnect.

Note: AS/400 Model 400 systems do not have external system buses, and, therefore, cannot be part of an OptiConnect cluster.

An OptiConnect cluster can consist of up to 32 systems, including AS/400 CISC systems. It supports interoperability between OS/400 versions so that each system can operate at different release levels from V3R1 and later.

D.1.2 Satellite and Hub Systems

The systems in an OptiConnect cluster share a common external optical system bus located in an expansion tower or frame. The system providing the shared system bus is called the hub system. Each system that plugs into this shared bus with an OptiConnect Bus Receiver card is called a satellite system. Each satellite system dedicates one of its external system buses that connects to the receiver card in the hub system's expansion tower or rack.

D.2 Link and Path Redundancy

The term OptiConnect *link* refers to the fiber optic cable connection between systems in the OptiConnect cluster. The term *path* refers to the logical software established connection between two OptiConnect systems.

There are two levels of redundancy available in an OptiConnect cluster:

1. Link redundancy
2. Path redundancy

D.2.1 Link Redundancy

Link redundancy is an optical bus hardware feature. Each optical link processor card has two external optical bus ports. The even numbered bus is the top connector and the odd numbered bus is the bottom connector. Both buses connect to bus receiver cards. When the bus receiver card has an available port, the bus receiver cards are linked with an additional fiber optic cable. When the primary cable for the bus fails, the Optical Link Processor (OLP) detects the failure and routes subsequent bus traffic across the other bus, passing through the first bus receiver card in the loop. This feature is available on CISC, OptiConnect, and RISC bus Models 530 and 53S.

Any two systems attached to the hub system shared bus can establish a path between them, including paths to the hub system itself. You can establish path redundancy by configuring two hub systems in the OptiConnect cluster. Each satellite uses two buses to connect with two hub systems. OptiConnect software detects the two logical paths between the two systems and uses both paths for data flow. If a path failure occurs, the remaining path picks up all of the communication traffic.




Figure 115. Link Redundancy

D.2.2 Path Redundancy

The second level of redundancy for the OptiConnect cluster is path redundancy. The OS/400 infrastructure for any system determines the logical path to another system. It does this by designating which system bus each of the systems that form the path uses.

The link between any two satellite systems does not depend on the hub system bus. The two systems use the bus, but the hub system is not involved.

Link redundancy is determined by the system models. For OptiConnect clusters, link redundancy is always provided when the extra fiber optic cable is installed. For path redundancy, an extra set of OptiConnect receiver cards and an extra expansion tower or frame are required, along with another set of cables.

D.3 Hardware Overview

All RISC Advanced Series AS/400 systems with external SPD buses support OptiConnect at a speed of 1063 Mbps. They also support the slower 266 Mbps OptiConnect hardware, which allows the satellite system to be located up to two kilometers away.

D.3.1 OptiConnect Adapter Cards and Connecting to the Network

RISC AS/400 systems use bus adapters that support 2686 OLPs or 2688 OLPs. Additional OLP cards for OptiConnect are ordered as an RPQ. The following diagram explains which adapter cards to install in each RISC model.

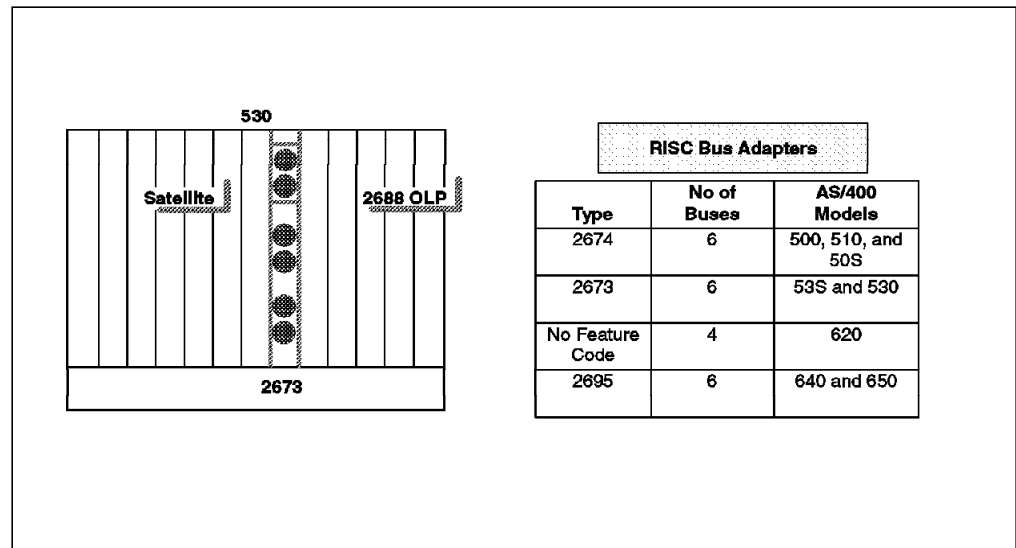


Figure 116. RISC Bus Adapter Hardware

In an OptiConnect network, it is helpful to understand how the AS/400 bus is used. First, an OptiConnect solution always consists of at least two AS/400 systems. One of these AS/400 systems is designated as the hub and at least one other system as a satellite.

The hardware to create a hub system for an OptiConnect cluster consists of a standard System Unit Expansion tower or frame with a Central Processing Unit (CPU) and a system unit.

All system expansion buses on the AS/400 connect to the system with optical cables. Systems can be connected together in the same manner. The possible cable connections, starting from the bus card, include those identified in Table 68 on page 390.

<i>Table 68. OptiConnect Cable Connections</i>					
AS/400 Model	Bus Card	Bus Pairing	Expansion Card/Tower	Description	Data Rate (Mbps)
RISC 53X 6XX SXX	2688	Yes	2682/5072	I/O Tower	1063 Mbps
			2682/5073	I/O Tower	
	2686	Yes	2685/	OptiConnect Card	266 Mbps
			2684/5044	I/O Frame	
500 50S 510	2686	No	2683/	OptiConnect Card	266 Mbps
		Yes	2680/5070	I/O Tower	
	2688	Yes	2684/5044	I/O Frame	
CISC Fxx, 3xx	Note*	Yes	2685	OptiConnect Card	1063 Mbps
			2632/5042	I/O Frame, two buses	220 Mbps
			2670/5042	I/O Frame, one bus	
			2670/5061	DASD Tower	
			2670/5062	I/O Tower	
			2669	OptiConnect to System B	

Note: The Bus Card can be of any of the following cards: Feature Code 2529, 2547, 2548, 2550, 2565, 2566, or 259F.

Appendix E. High Availability Solutions

The AS/400 development and manufacturing teams continue to improve the AS/400 system. For customers needing better than 99.9% system availability, AS/400 clusters are available. Cluster solutions connect multiple AS/400 systems together with various interconnect fabrics, including high-speed optical fiber, to offer a solution that can deliver up to 99.99% system availability.

Combining these clusters with software from AS/400 high-availability business partners – such as those described in this appendix – improves the availability of a single AS/400 by replicating business data to one or more AS/400 systems. This combination can provide a disaster recovery solution.

In recent AS/400 hardware and operating system releases, many improvements were made to the availability and recovery functions. These enhancements help AS/400 end-users achieve greater availability of application and data services on a single system. For planned or unplanned outages, clustering and system mirroring offer the most effective solution. IBM business partners that provide high systems availability tools continue to complement IBM's availability offerings with clustering and system mirroring solutions.

OS/400 functions provide the foundation for the High Availability Solutions described in this appendix. The OS/400 V4R2 remote journal enhancement, in particular, enhances these solutions by enabling functions below the machine interface (MI) level. These functions were previously coded into application programs.

For more information about the remote journal function, see Chapter 17, "Using Remote Journals to Improve Availability and Recovery" on page 327, in this redbook.

E.1 A High Availability Customer Scenario

A Danish customer with 3000 users on a large AS/400 system had difficulty finding time to complete tasks that required a dedicated system for maintenance. These tasks included such operations as performing nightly backups, installing new releases, and updating the hardware.

One reason this challenge occurred is because the customer's AS/400 system was serving all of their retail shops across Scandinavia. As these shops extended their hours, there was less and less time for planned system outages.

They solved their problem by installing mirroring software on two AS/400 systems. This solution made it easy for the shops to expand their hours and improve sales without losing system availability or sacrificing system maintenance.

To increase availability, the customer bought a second AS/400 system and connected the two machines with OptiConnect/400. Next, they installed mirroring software, and mirrored everything on the production machine to the backup machine. In the event of a planned or an unplanned system outage, the system's users could switch to the backup machine in minutes.

E.2 When to Consider a High Availability Solution

When considering if a high availability solution is right for you, ask yourself these questions:

- Will we benefit from using synchronized distributed databases?
- Do our users need access to the AS/400 system 24 hours a day, 365 days a year?
- Do our users operate in different time zones?
- Is there enough time for nightly backups, scheduled maintenance, or installing new releases?
- If our telephone sales application is not always up and running, will we lose our customers to the competition?
- Is there a single point of failure for any data center?
- Can we avoid the loss of data or access to the system in the event of a disaster or sabotage?
- When the production machine is overloaded, can we move some users to a different machine for read-only jobs?

A high availability solution can benefit any or each of these situations.

E.2.1 What a High Availability Solution Is

High (or continuous) availability systems usually include an alternate system or CPU that mirrors some of the activity of the production system, and a fast communications link. These systems also include replication or mirroring software and enough DASD to handle the volume of data for a reasonable recovery time as shown in the following diagram.




Figure 117. The Basics in a High Availability Solution

This appendix demonstrates that High Availability Solutions exist. It also highlights some of the functions that are relevant to AS/400 high-availability considerations.

In this appendix, we outline four High Availability Solutions. Three are offered by IBM business partners:

- DataMirror Corporation
- Lakeview Technology
- Vision Solutions, Inc.

IBM offers a solution that fulfills some of the requirements of an HSA solution:

- DataPropagator Relational/400

The DataPropagator Relational/400 product was not designed as a high availability solution. In some cases, it can cover the needs for data availability, as discussed in the next section.

The following table outlines some of the requirements of a HSA solution to help simplify your investigation of High Availability Solutions.

<i>Table 69 (Page 1 of 2). Requirements of a High Availability Solution</i>				
Features	Solution A	Solution B	Solution C	DataPropagator
24x7 availability				No
Eliminate downtime for backup and maintenance				No
Replication of database				Yes
Replication of other objects				No
Data replication to non-AS/400 systems				Yes
Handle unplanned outages				No
Automatically switch users to a target system				No
Workload distribution				Yes
Error recovery				Yes
Distribution to multiple AS/400 systems				Yes
Commitment control support				
Sync checks				Yes
Filtering of mirrored objects				Yes (DB only)
Execute remote commands				No
OptiConnect support				Yes

Table 69 (Page 2 of 2). Requirements of a High Availability Solution				
Features	Solution A	Solution B	Solution C	DataPropagator
Utilize Remote Journals				Yes

Note: Some of the software vendors mentioned in this appendix may have products with functions that are not directly related to the high availability issues on the AS/400 system. To learn more about these products, visit these vendors on the World Wide Web. You can locate their URL address at the end of the section that describes their solution.

E.3 DataMirror

DataMirror Corporation, an IBM business partner, has products that address a number of issues such as data warehousing, data and workload distribution, and high availability. DataMirror’s products run on IBM and non-IBM platforms.

DataMirror’s High Availability Suite uses high performance replication to ensure reliable and secure delivery of data to backup sites. In the event of planned or unplanned outages, the suite ensures data integrity and continuous business operations. To avoid transmission of redundant data, only changes to the data are sent to the backup system. This allows resources to be more available for production work. After an outage is resolved, systems can be resynchronized while they are active.

Figure 118 illustrates the components of DataMirror’s High Availability Suite.

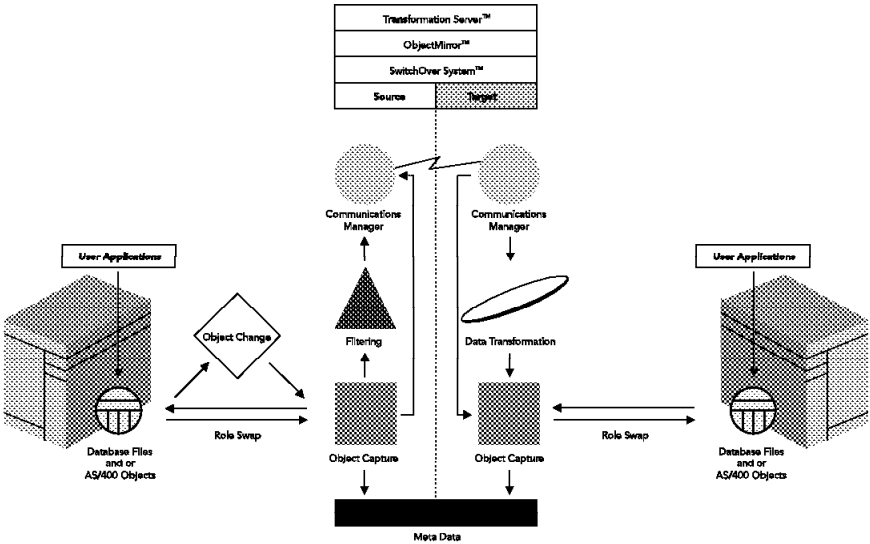


Figure 118. DataMirror High Availability Suite

DataMirror’s High Availability Suite contains three components:

- DataMirror High Availability (HA) Data
- ObjectMirror
- SwitchOver System

The following section highlights each component.

E.3.1 DataMirror HA Data

DataMirror HA Data mirrors data between AS/400 production systems and fail-over machines for backup, recovery, high systems availability, and clustering.

A user can replicate entire databases or individual files on a predetermined schedule in real-time or on a net change basis. They can refresh the backup machine nightly or weekly as required. Or they can use DataMirror HA Data to replicate changes to databases in real-time so that up-to-the-minute data is available during a scheduled downtime or disaster.

DataMirror HA Data software is a no-programming-required solution. Users simply install the software, select which data to replicate to the backup system, determine a data replication method (scheduled refresh or real-time), and begin replication. At the end of a system failure, fault tolerant resynchronization can occur without taking systems offline.

DataMirror HA Data supports various high availability options including workload balancing, 7 X 24 hour operations availability, and critical data backup. Combined with Data Mirror's ObjectMirror software and SwitchOver System, a full spectrum of high availability options is possible.

E.3.2 ObjectMirror

ObjectMirror enables critical application and full system redundancy to ensure access to both critical data and the applications that generate and provide use of the data.

ObjectMirror supports real-time object mirroring from a source AS/400 system to one or more target systems. It provides continuous mirroring, as well as an on-demand full refresh of AS/400 objects, which are grouped by choice of replication frequency and priority.

ObjectMirror's features include:

- Grouping by choice to mirror like-type objects based on frequency or priority
- Continuous real-time mirroring of AS/400 objects
- Intelligent replication for guaranteed delivery to backup systems even during a system or communication failure
- Object refreshment on a full-refresh basis as needed
- Fast, easy setup including an automatic registration of objects
- Ability to send an object or group of objects immediately without going through product setup routines

E.3.3 SwitchOver System

The SwitchOver System operates on both the primary and backup AS/400 systems to monitor communications or system failures. During a failure, the SwitchOver System initiates a logical role switch of the primary and backup AS/400 systems either immediately or on a delayed basis.

A Decision Control Matrix in the SwitchOver System allows multiple line monitoring, detailed message logging, automated notification, and user-exit processing at various points during the switching process. An A/B switch (as shown in Figure 119 on page 396) lets the user automatically switch users and hardware peripherals, such as twinax terminals, printers, and remote controllers.

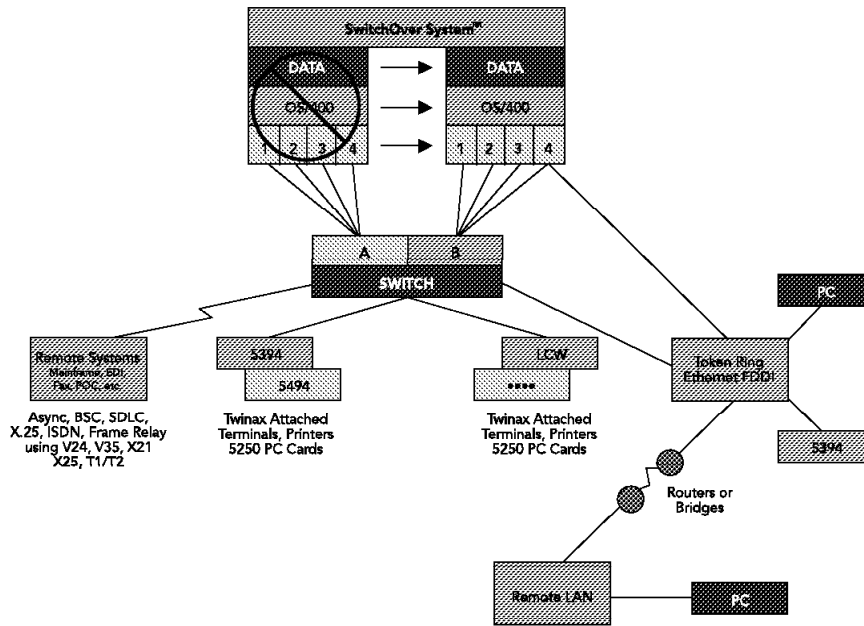


Figure 119. DataMirror SwitchOver System

E.3.4 OptiConnect and DataMirror

DataMirror's High Availability Suite supports SNA running over OptiConnect between AS/400 systems. After OptiConnect is installed on both source and target AS/400 systems, the user needs to create controllers and device descriptions.

Once controllers and devices are varied on, the user simply specifies the device name and remote location used in the DataMirror HA Data or Object Mirror target definition. Files or objects that are specified can then be defined, assigned to the target system, and replicated.

E.3.5 Remote Journals and DataMirror

The DataMirror High Availability (HA) Suite is capable of using the IBM remote journal function in OS/400 V4R2. The architecture of the DataMirror HA Suite allows the location of the journal receivers to be independent from where the production (source) or fail-over (target) databases reside. Therefore, journal receivers can be located on the same AS/400 system as the fail-over database, allowing the use of DataMirror's "intra-system" replication to support remote journals.

Customers can invoke remote journal support in new implementations. Or, the existing setup can be modified if remote journal support was not originally planned.

E.3.6 More Information about DataMirror

To learn more information about DataMirror products, visit DataMirror on the Internet at:

<http://www.datamirror.com>

E.4 IBM and High Availability

IBM's contribution to AS/400 High Availability Solutions includes the IBM DataPropagator Relational Capture and Apply for AS/400 product. From this point forward, we refer to this product as DataPropagator/400. This section describes IBM's package and its benefits as a minimal High Availability Solution.

E.4.1 IBM DataPropagator Relational Capture and Apply for AS/400

DataPropagator/400 is a state-of-the-art data replication tool. Data replication is necessary when:

- Supplying consistent real-time reference information across an enterprise
- Bringing real-time information closer to the business units that require access to insulate users from failures elsewhere on the network
- Reducing network traffic or the reliance on a central system
- On-demand access disrupts production or response
- Migrating systems and designing a transition plan to move the data while keeping the systems in sync
- Deploying a data warehouse with an automated movement of data
- Current disaster plan strategies do not adequately account for site-failure recovery

DataPropagator/400 is not a *total* High Availability Solution because it only replicates databases. It does not replicate all the objects that must be mirrored for a true High Availability Solution in a dynamic environment.

However, consider DataPropagator/400 for availability functions in a stable environment where the following criteria can be met:

- Only the database changes during normal production on the AS/400.
- Such objects as user profiles, authorities, and other nondatabase objects are saved regularly on the source system and restored on the target system when changed.

In other words, in a stable environment where only the database changes, replicating the database to a backup system and transferring users manually to this system may be a sufficient availability and recovery plan.

See Figure 120 on page 398, which illustrates DataPropagator/400 as a data replicator tool.

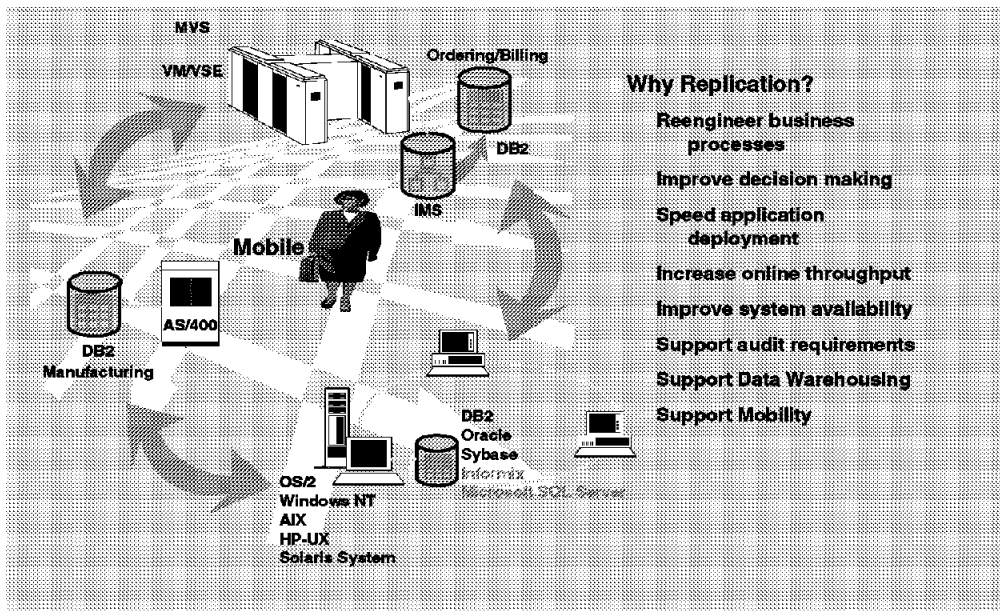


Figure 120. Usage of DataPropagator

E.4.2 DataPropagator/400 Description

The IBM DataPropagator Relational Capture and Apply for AS/400 automatically copies data within and between IBM DB2 platforms to make data available on the system when it is needed. The IBM DataJoiner product can be used in addition to the DataPropagator product to provide replication to several non-IBM databases. Immediate access to current and consistent data reduces the time necessary for analysis and decision making.

DataPropagator Relational/400 allows the user to update copied data, maintain historical change information, and control copy impact on system resources. Copying may involve transferring the entire contents of a user table (a full refresh) or only the changes since the last copy (an update). The user can also copy a subset of a table by selecting the columns they want to copy.

Making copies of database data (snapshots) is a solution to the problem of remote data access and availability. Copied data requires varying levels of synchronization with production data depending on how the data is used.

Copying data may even be desirable within the same database. If excessive contention occurs for data access in the master database, copying the data offloads some of the burden from the master database.

By copying data, users can get information without impacting their production applications. It also removes any dependency on the performance of remote data access and the availability of communication links.

DataPropagator Relational/400 highlights include:

- An automatic copy of databases
- Full support for SQL (enabling summaries, derived data, and subsetted copies)
- During a system or network outage, the product restarts automatically from the point where it stopped. If this is not possible, a complete refresh of the copies can be performed if the administrators allow it. And, for example, if one of the components fails, the product can determine that there is a break in sequence of the data being copied. In this case, DataPropagator restarts the copy from scratch.
- Open architecture to enable new applications
- DataPropagator/400 commands that support AS/400 system definitions
- Full usage of remote journal support in V4R2

E.4.3 DataPropagator/400 Configuration

In the database network, the user needs to assign their systems one or more roles when configuring the DataPropagator environment, including:

- Control server

This system contains all the information on the registered tables, the snapshot definitions (the kind of data we want to copy and how to copy it), the ownership of the copies, and the captures in reference to registrars and subscribers.

- Data server

Contains the source data tables.

- Copy server

This is the target system.

Depending on the structure of the company, the platforms involved and the customer's preferences, a system in the network can play one or more of these three roles. DataPropagator, for instance, works powerfully on a single AS/400 system, which, at the same time, serves as Control, Copy, and Data Server.

E.4.4 Data Replication Process

With DataPropagator, there are two steps to the data replication process:

- The *Capture* process—for reading the data
- The *Apply* process—for applying updated data

The following two graphics illustrate these processes.

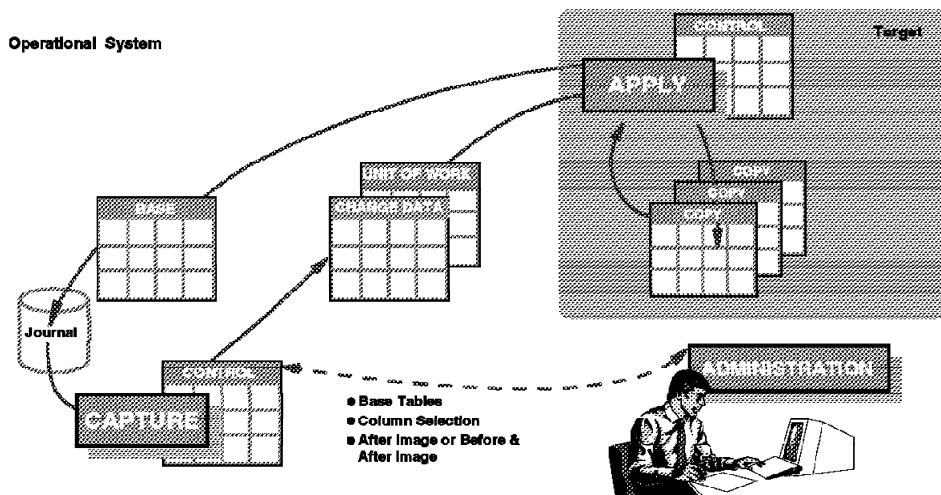


Figure 121. The DataPropagator Capture Process

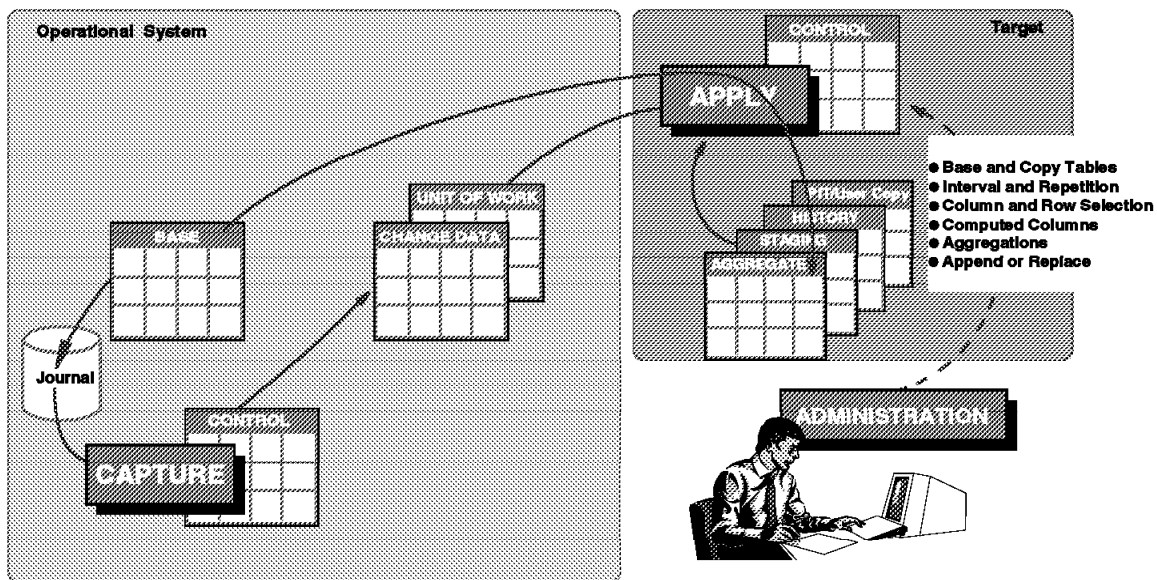


Figure 122. The DataPropagator Apply Process

Enhancements in V4R2 include:

- Support for the remote journal function to offload the source CPU
- Automated deletion of journal receivers
- Replication over native TCP/IP

- Multi-vendor replication with DataJoiner (replication to and from Oracle, Sybase, Informix, and Microsoft SQL Server databases)
- Integration with the Lotus Notes databases

E.4.5 OptiConnect and DataPropagator/400

DataPropagator is based on a distributed relational database architecture (DRDA) and is independent of any communications protocol. Therefore, it uses OptiConnect and any other media without additional configuration.

E.4.6 Remote Journals and DataPropagator/400

DataPropagator/400 takes advantage of the operating system's remote journal function. With remote journals, the capture process is run at the remote journal location to offload the capture process overhead from the production system.

The apply process does not need to connect to the production system for differential refresh because the DataPropagator/400 staging tables reside locally, and not on the production system. In addition, because the DataPropagator/400 product is installed only on the system that is journaled remotely, the production system no longer needs a copy of DataPropagator/400.

E.4.7 DataPropagator/400 Implementation

DataPropagator/400 is most beneficial for replicating data to update remote databases. One real-life example of this is a customer in Denmark who had a central AS/400 system and stored all production data, pricing information, and a customer database on it. From this central machine, data was distributed to sales offices in Austria, Germany, Norway, and Holland, each of which operated either small AS/400 systems or OS/2 PCs. Each sales office received a subset of the data that was relevant to their particular office.

E.4.8 More Information about DataPropagator

For more information about IBM DataPropagator solutions, refer to *DataPropagator Relational Guide*, SC26-3399, *DataPropagator Relational Capture and Apply/400*, SC41-5346-01, and visit the IBM Internet site at:

<http://www.software.ibm.com/data/dbtools/datarepl.html>

E.5 Lakeview Technology

Lakeview Technology, an IBM business partner, offers a number of products to use in an AS/400 high availability environment. Their high availability suite contains five components:

- MIMIX/400
- MIMIX/Object
- MIMIX/Switch
- MIMIX/Monitor
- MIMIX/Promoter

The following sections highlight each of these components.

E.5.1 MIMIX/400

MIMIX/400 is the lead module in Lakeview Technology's MIMIX high availability management software suite for the IBM AS/400 system. It creates and maintains one or more exact copies of a DB2/400 database by replicating application transactions as they occur. The AS/400 system pushes the transaction data to one or more companion AS/400 systems. That way, a viable system with up-to-date information is always available when planned maintenance or unplanned disasters bring down the primary system. MIMIX/400 also supports intra-system database replication.

The following graphic shows the basic principles of the MIMIX/400.

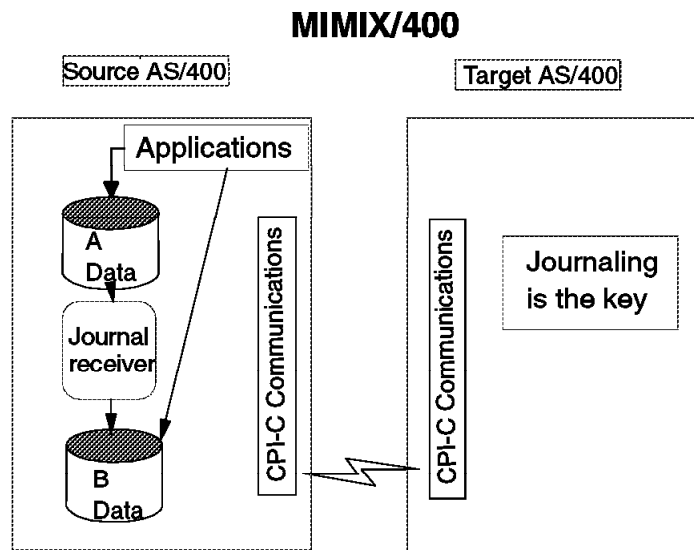


Figure 123. The Basics Principles of Mimix/400

The key functions of MIMIX/400 are:

- Send
- Receive
- Apply
- Synchronize
- Switch

The *Send* function scrapes the source system's journal and sends the data to one or more target systems. This function offers these characteristics:

- Is written in ILE/C for high performance
- Use CPI-C to provide a low-level generic interface that keeps CPU overhead to a minimum
- Supports filtering to eliminate files from MIMIX copies, and to optimize communication throughput, auxiliary storage usage, and performance on the target system
- Generates performance stamps, which continue throughout the replication cycle for a historic view of performance bottlenecks

The *Receive* function collects transactions from the Send function. The Receive function stores and manages the transactions on the backup system for processing using the Apply function. The Receive function offers these features:

- A temporary staging step where transactions are pushed off the sending system as soon as possible to eliminate the load from the source system's CPU
- Fast performance because it is written in ILE/C
- Variable length log-space entries to make the most of available CPU and DASD resources
- Filtering capabilities for greater capacity exists on the target system to boost performance on the Send side

The *Apply* function reads all transactions and updates the duplicate databases on the target systems accordingly. The Apply function supports these features:

- Offers a file or member control feature to manage file name aliases and define files to the member level (files are locked during the Apply process for maximum configuration flexibility and to prevent files from being unsynchronized)
- Opens up to 9999 files simultaneously within a MIMIX Apply session
- Supports record lengths up to 32K in size
- Manages DB2/400 commitment control boundaries during the Apply and Switch processes
- Uses a log process to protect against data loss during a source system outage
- Includes a graphical status report of source and target system activity. It displays the report in an easy-to-read format for operators to quickly identify MIMIX operating environment issues.

The *Synchronize* function verifies that the target system has recorded exact copies of the source system data. The *Synchronize* function supports these features:

- Offers keyed synchronization to keep target and source databases in synchronization with the unique "key" field in each record
- Provides support tools to analyze and correct file synchronization errors by record

The *Switch* function prepares target systems for access by users during a source system outage. The Switch function performs the following tasks:

- Defines systems, journals, fields, and data areas. The Send, Receive, and Apply sessions are linked into a logical unit called a data group.
- Uses a data group manager to reverse the direction of all MIMIX/400 transmissions during an outage
- Offers a journal analysis tool to identify transactions that may be incomplete after an outage

E.5.2 MIMIX/Object

The MIMIX/Object component creates and maintains duplicate images of critical AS/400 objects. Each time a user profile, device description, application program, data area, data queue, spool file, PC file, image file, or other critical object is added, changed, moved, renamed or deleted on an AS/400 production system, MIMIX/Object duplicates the operation on one or more backup systems.

The key elements of MIMIX/Object include:

- Audit Journal Reader
- Distribution Reader
- Send Network Object

Audit Journal Reader scrapes the source system's security audit journal for object operations and passes them to the distribution reader. The features of the Audit Journal Reader include:

- Management of objects within a library by object and type; document library objects (DLOs) by folder path, document name, and owner; and integrated file system objects by directory path and object name.
- Management of spooled-file queues based on their delivery destination
- Explicit, generic (by name), and comprehensive (all) identification of library, object, DLO, and integrated file system names
- An "include" and "exclude" flag for added naming precision
- Integrated file system control to accommodate hierarchical directories, support long names, and provide additional support for byte stream files

The *Distribution Reader* sends, receives, confirms, retries, and logs objects to history and message queues. The features of the Distribution Reader include:

- Multi-thread asynchronous job support to efficiently handle high volumes of object operations
- A load-leveling journal monitor to automatically detect a large backlog for greater parallelism in handling requests
- A history log to monitor successful distribution requests; offering reports by user, job, and date; and effective use of time for improving security control and management analysis
- A failed request queue to provide error information, and to delete and retry options for ongoing object integrity and easy object resolution
- An automatic retry feature to resubmit requests when objects are in use by another application until the object becomes available
- Automatic management of journal receivers, history logs, and transaction logs to minimize the use of auxiliary storage

The *Send Network Object* relies on the Audit Journal Reader, which interactively saves and restores any object from one system to another. It offers:

- Simplified generic distribution of objects manually or automatically through batch processing

E.5.3 MIMIX/Switch

The MIMIX/Switch component detects system outages and initiates the MIMIX recovery process. It automatically switches users to an available system where they can continue working without losing information or productivity.

The key elements of MIMIX/Switch include:

- Logical Switch
- Physical Switch
- Communications Monitor

The *Logical Switch* controls the physical switch, communication and device descriptions, network attributes, APPC/APPN configurations, TCP/IP attributes, and timing of the communication switchover. The features of the Logical Switch include:

- User exits to insert user-specified routines almost anywhere in the command stream to customize the switching process
- A message logging feature to send status messages to multiple queues and logs for ensuring the visibility of critical information

The *Physical Switch* automatically and directly communicates with the gang switch controller to create a switch over. The features of the *Physical Switch* include:

- A custom interface to the gang switch controller to switch communication lines directly
- An operator interface to facilitate manual control over the gang switch controller
- Remote support to initiate a switch through the gang switch controller from a distance
- Interface support of: twinax, coax, RJ11, RS232, V.24, v.35., X.21, DB9 or other devices that the user can plug into a gang switch

The *Communications Monitor* tracks the configuration object status to aid in automating retry and recovery. An automatic verification loop ensures that MIMIX/Switch only moves users to a backup system when a genuine source system outage occurs.

E.5.4 MIMIX/Monitor

The MIMIX/Monitor component combines a command center for the administration of monitor programs and a library of plug-in monitors so the user can track, manage, and report on AS/400 processes.

MIMIX/Monitor regulates the system 24 hours a day. It presents all monitor programs on a single screen with a uniform set of commands, minimizing the time and effort required to insert or remove monitors or change their parameters. The MIMIX/Monitor also accepts other data monitoring tools created by customers and third-party companies into its interface.

The user can set the programs included with MIMIX/Monitor to run immediately, continually at scheduled intervals, or after a particular event (for example, a communications restart).

MIMIX/Monitor includes prepackaged monitor programs that the user can install to check the levels in an uninterruptible power supply (UPS) backup system, or to evaluate the relationship of MIMIX to the application environment.

E.5.5 MIMIX/Promoter

The MIMIX/Promoter component helps organizations maintain continuous operations while carrying out database reorganizations and application upgrades, including year 2000 date format changes. It uses data transfer technology for revising and moving files to production without seriously affecting business operations.

MIMIX/Promoter builds copies of database files record-by-record, working behind the scenes while users maintain read-and-write access to their applications and data. It allows the user to fill the new file with data, change field and record lengths, and at the same time, keep the original file online.

After copying is complete, MIMIX/Promoter moves the new files into production in a matter of moments. This is the only time when the application must be taken off line.

Implementing an upgrade also requires promoting such nondatabase objects as programs and display files. To handle these changes, many organizations use change management tools, some of which can be integrated with MIMIX/Promoter's data transfer techniques.

E.5.6 OptiConnect and MIMIX

MIMIX/400 integrates OptiConnect for OS/400 support for IBM's high-speed communication link, without requiring separate modules. The combination of MIMIX and OptiConnect provides a horizontal growth solution for interactive applications that are no longer contained on a single machine. OptiConnect delivers sufficient throughput for client/server-style database sharing among AS/400 systems within a data center for corporate use. MIMIX/400 complements the strategy by making AS/400 server data continuously available to all clients.

E.5.7 More Information About Lakeview Technology

For more information about Lakeview Technology's complete product line, visit Lakeview Technology on the Internet at:

<http://www.lakeviewtech.com>

E.6 Vision Solutions, Inc.

Vision Solutions, Inc. products operate on two or more AS/400 systems in a network and use mirroring techniques. This ensures that databases, applications, user profiles, and other objects are automatically updated on the backup machines. In case of a system failure, end users and network connections are automatically transferred to a predefined backup system. The Vision products automatically activate the backup system (perform a role swap) without any operator intervention.

With this solution, two or more AS/400 systems can share the workload. For example, it can direct end-user queries that do not update databases to the backup system. Other benefits of this solution include dedicated system

maintenance projects. The user can temporarily move their operations to the backup machine and upgrade or change the primary machine.

This High Availability Solution offers an easy and structured way to keep AS/400 business applications and data available 24 hours a day, 7 days a week.

The Vision Solutions, Inc. High Availability Solution, called Vision Suite, includes three components:

- Object Mirroring System (OMS/400)
- Object Distribution System (ODS/400)
- System Availability Monitor (SAM/400)

The following sections highlight each component.

E.6.1 OMS/400—Object Mirroring System

The Object Mirroring System (OMS/400) automatically maintains duplicate databases across two or more AS/400 systems.

Figure 124 illustrates the OMS/400 system. This system uses journals and a communication link between the source and target systems.

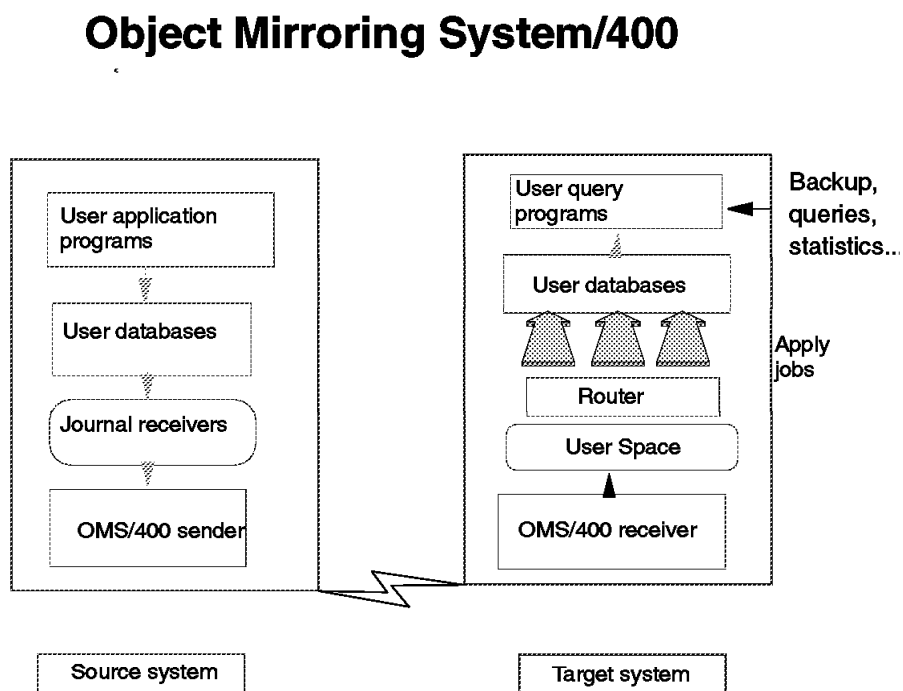


Figure 124. The Object Mirroring System/400 Mirroring Process

The features of the OMS/400 component include:

- Automatic repair of such abnormal conditions as communication, synchronization, or system failure recoveries
- Synchronization of enterprise-wide data by simulcasting data from a source system to more than 9 000 target destinations
- User space technology that streamlines the replication process

- An optional ongoing validity check to ensure data integrity
- Automatic restart after any system termination
- Automatic filtration of unwanted entries, such as opens and closes
- The power to operate programs or commands from a remote system
- The ability to dynamically capture data and object changes on the source system and copy them to the target system without custom commands or recompiles
- The option to create an unlimited number of prioritized AS/400 links between systems
- Total data protection by writing download transactions to tape
- Support of RPG/400 for user presentation, and ILE/C for system access, data transmission, and process application
- Full support of the IBM OptiConnect/400 system
- Global journal management when a fiber optic bus-to-bus connection is available
- The use of CPI-C to increase speed of data distribution using minimal CPU resources

E.6.2 ODS/400—Object Distribution System

The Object Distribution System (ODS/400) provides automatic distribution of application software, authority changes, folders and documents, user-profile changes, and system values. It also distributes subsystem descriptions, job descriptions, logical files, and output queue and job queue descriptions.

ODS/400 is a partner to the OMS/400 system, and provides companies with full system redundancy. It automatically distributes application software changes, system configurations, folders and documents, and user profiles throughout a network of AS/400 computers.

ODS/400 supports multi-directional and network environments in centralized or remote locations. For maximum throughput, ODS/400 takes advantage of bi-directional communications protocol and uses extensive filters.

E.6.3 SAM/400—System Availability Monitor

The System Availability Monitor (SAM/400) can switch users from a failed primary system to their designated secondary system without operator intervention.

SAM/400 works in conjunction with OMS/400 and ODS/400, continuously monitoring the source system. In the event of a failure, SAM/400 automatically redirects users to the target system, virtually eliminating downtime. High-speed communications links, optional electronic switching hardware, and SAM/400 work together to switch users to a recovery system in only a few minutes.

SAM/400 offers:

- Continuous monitoring of all mirrored systems for operational status and ongoing availability
- A fully programmable response to react automatically during a system failure, which reduces the need to depend on uninformed or untrained staff

- The ability to immediately and safely switch to the target system, which contains an exact duplicate of the source objects and data during a source system failure (unattended systems are automatically protected 24 hours a day, 7 days a week)
- User-defined access to the target system based on a specific user class or customized access levels

The SAM/400 component offers:

- Up to ten alternate communication links for monitoring from the target system to the source system
- Automatic initiation of user-defined actions when a primary system failure occurs
- Exit programs to allow the operator to customize recovery and operations for all network protocols and implementations

Figure 125 illustrates the SAM/400 monitoring process.

System Availability Monitor/400

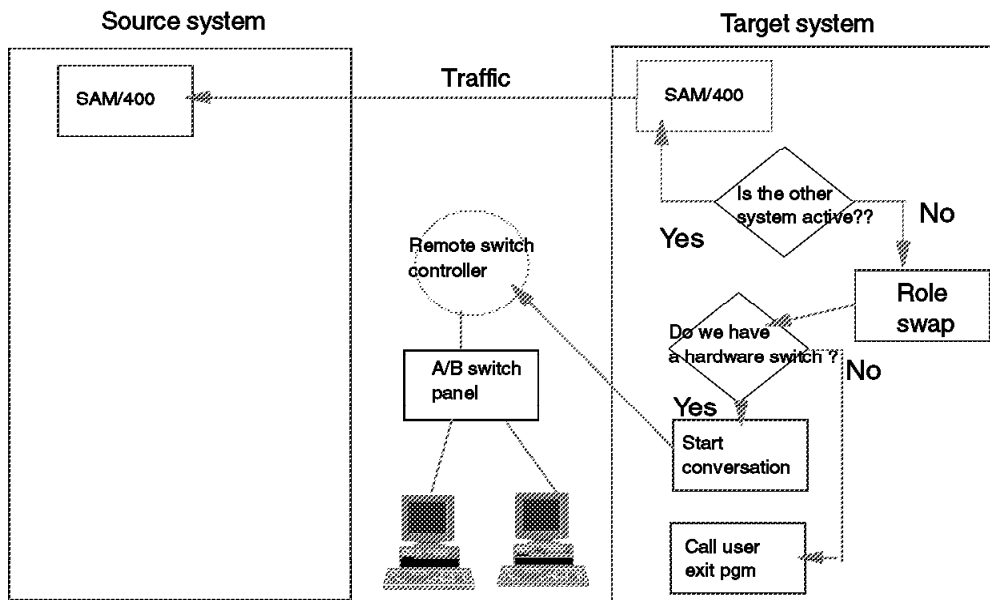


Figure 125. SAM/400 Structure

Users are allowed to access applications at the “End” point.

E.6.4 High Availability Services/400

High Availability Services/400 (HAS/400) consists of software and services. The HAS/400 solution is comprised of:

- Analysis of the customer’s environment in terms of system availability needs and expectations, critical business applications, databases, and workload distribution capabilities

- An implementation plan written in terms of solution design and the required resources for its deployment
- Installation and configuration of the software products and the required hardware
- Education for the customer's staff on operational procedures
- Solution implementation test and validation
- Software from Vision Solutions, Inc., as previously described

E.6.5 More Information About Vision Solutions, Inc.

For more information about Vision Solutions, Inc. products, visit Vision Solutions on the Internet at:

<http://www.visionsolutions.com>

Appendix F. Special Notices

This publication is intended to help the system administrator find information about availability, backup, and recovery options, enhancements, and service offerings. The information in this publication is not intended as the specification of any programming interfaces that are provided by high availability vendors, such as DataMirror Corporation, Lakeview Technology, and Vision Solutions, Inc. See the PUBLICATIONS section of the IBM Programming Announcement for more information about what publications are considered to be product documentation.

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.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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 Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. 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.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other

operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

AnyNet®	Application System/400®
APPN®	AS/400 PerformanceEdge
AS/400®	Business Partner
Client Access	Client Access/400
Client Series	COBOL/400®
DATABASE 2	DATABASE 2 OS/400
DataJoiner®	DataPropagator
DB2®	DB2® Connect
DB2® Client Application Enablers	DB2® Universal Database
DB2® Universal Server	Distributed Relational Database Architecture
DProp	IBM Business Partner (logo)
IBM®	IBMLink
Information Assistant	Integrated Language Environment®
Magstar	Object Connection
OfficeVision®	OfficeVision/400
Operating System/2®	Operating System/400®
POWER Architecture	RETAIN®
RPG/400®	RS/6000
SystemView®	Workplace

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

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

Appendix G. Related Publications

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

G.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 417.

- *Fire in the Computer Room, What Now? Disaster Recovery, Planning for Business Survival*, SG24-4211
- *AS/400 System Availability and Recovery for V2R2*, GG24-3912
- *AS/400 Advanced Series System Builder V4R2*, SG24-2155
- *AS/400 Advanced Series System Handbook V4R2*, GA19-5486-16
- *Speak the Right Language with Your AS/400 System*, SG24-2154
- *AS/400 Performance Management V3R6/V3R7*, SG24-4735
- *AS/400 Implementing Windows NT on the Integrated PC Server*, SG24-2164
- *DB2/400 Advanced Database Functions*, SG24-4249
- *Database Parallelism on the AS/400*, SG24-4826
- *So You Want to Estimate the Value of Availability*, GG22-9318
- *AS/400 E-Commerce: Internet Connection Servers*, SG24-2150

G.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

G.3 Other Publications

These publications are also relevant as further information sources:

- *AS/400 Security — Basic V4R1*, SC41-5301
- *Security — Reference*, SC41-5302
- *OptiMover for OS/400 PRPQ V3R6*, SC41-0626
- *WRKASP Utility User's Guide PRPQ V3R6*, SC41-0652

- *WRKASP Utility User's Guide PRPQ V4*, SC41-0675
- *Tips and Tools for Securing Your AS/400 V4R2*, SC41-5300
- *AS/400 Licensed Internal Code Diagnostic Aids — Volume 1*, LY44-5900 (available to IBM licensed customers only)
- *AS/400 Licensed Internal Code Diagnostic Aids — Volume 2*, LY44-5901 (available to IBM licensed customers only)
- *Performance Tools/400*, SC41-5340
- *Backup Recovery and Media Services V3R7*, SC41-4345
- *Job Scheduler for OS/400 V3R6*, SC41-4324
- *Managed System Services/400 Use V3R7*, SC41-3323
- *Report/Data Archive and Retrieval System Installation and User's Guide V3R2*, SC41-3325
- *System Manager Use V4R2*, SC41-5321
- *SystemView for OS/400: Up and Running V3R7*, SC41-4330
- *TME 10 NetFinity for AS/400 V3R7*, SC41-4331
- *ADSM for OS/400: Administrator's Guide*, SC35-0196
- *ADSM for OS/400: Administrator's Reference*, SC35-0197
- *OS/400 Advanced Series ILE RPG/400 Reference Summary*, SX09-1306
- *OS/400 Advanced Series ILE RPG/400 Programmer's Guide*, SC09-2074
- *OS/400 Advanced Series ILE RPG/400 Reference Version 3*, SC09-2077
- *AS/400 Languages:RPG/400 Reference*, SC09-1817
- *AS/400 Languages:RPG/400 User's Guide*, SC09-1816
- *Distributed Database Programming V4R2*, SC41-5702
- *DB2 for AS/400 Database Programming V4R2*, SC41-5701
- *DB2 for AS/400 SQL Programming V4R2*, SC41-5611
- *AS/400 Alerts Support V4R1*, SC41-5413
- *APPN Support V4R2*, SC41-5407
- *Communications Configuration V4R1*, SC41-5401
- *Communications Management V4R2*, SC41-5406
- *AS/400 LAN and Frame Relay Support V4R2*, SC41-5404
- *Remote Work Station Support V4R1*, SC41-5402
- *OptiConnect for OS/400 V3R7*, SC41-4414
- *OptiConnect for OS/400 V3R2*, SC41-3414
- *DB2 for AS/400 Database Programming V4R2*, SC41-5701
- *DB2 Multisystem for AS/400 V4R1*, SC41-5705
- *DataPropogator Relational Capture and Apply for OS/400 V4R2*, SC41-5346
- *ADSM: Using the Application Programming Interface V2R1*, SH26-4002
- *ADTS/400: Application Development Manager API Reference V3R2*, SC09-1809
- *Common Programming APIs Toolkit/400 Reference V3R7*, SC41-4802

- *System API Programming V4R1*, SC41-5800
- *System API Reference V4R2*, SC41-5801
- *CL Programming V4R2*, SC41-5721
- *CL Reference V4R2*, SC41-5722
- *System/400 International Application Development*, SC41-5603
- *National Language Support V4R2*, SC41-5101
- *AS/400 Physical Planning Reference V4R2*, SA41-5109
- *AS/400 Software Installation V4R2*, SC41-5120
- *TCP/IP Configuration and Reference*, SC41-5420
- *Firewall for AS/400 V4R1*, SC41-5424
- *Internet Connection Server and Internet Connection Secure Server for AS/400 Webmaster's Guide V4R2*, GC41-5434
- *Integrating AS/400 with Novell NetWare V4R1*, SC41-5124
- *Integration Services for the Integrated PC Server V4R1*, SC41-5123
- *OS/2 Warp Server for AS/400 Administration V4R1*, SC41-5423
- *AS/400 Client Access Host Servers V4R2*, SC41-5740
- *OS/400 Integration of Lotus Notes V4R1*, SC41-5431
- *Basic System Operation, Administration, and Problem Handling*, SC41-5206
- *AS/400 System Operation V3R6*, SC41-4203
- *AS/400 System Operation for New Users V3R1*, SC41-3200
- *Automated Tape Library Planning and Management V4R2*, SC41-5309
- *OS/400 Backup and Recovery V4R2*, SC41-5304
- *OS/400 Central Site Distribution V4R1*, SC41-5308
- *OS/400 Work Management V4R2*, SC41-5306
- *Distributed Database Programming V4R2*, SC41-5702
- *Data Management V4R1*, SC41-5710
- *OS/400 Distributed Data Management V4R1*, SC41-5307
- *Integrated File System Introduction V4R2*, SC41-5711

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at <http://www.redbooks.ibm.com/>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Redbooks Web Site on the World Wide Web**

<http://w3.itso.ibm.com/>

- **PUBORDER** — to order hardcopies in the United States

- **Tools Disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLCAT REDPRINT
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get BookManager BOOKs of redbooks, type the following command:

```
TOOLCAT REDBOOKS
```

To get lists of redbooks, type the following command:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
```

- **REDBOOKS Category on INEWS**

- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

In United States:
In Canada:
Outside North America:

IBMMAIL
usib6fpl at ibmmail
caibmbkz at ibmmail
dkibmbsh at ibmmail

Internet
usib6fpl@ibmmail.com
lmannix@vnet.ibm.com
bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)
Canada (toll free)

1-800-879-2755
1-800-IBM-4YOU

Outside North America
(+45) 4810-1320 - Danish
(+45) 4810-1420 - Dutch
(+45) 4810-1540 - English
(+45) 4810-1670 - Finnish
(+45) 4810-1220 - French

(long distance charges apply)
(+45) 4810-1020 - German
(+45) 4810-1620 - Italian
(+45) 4810-1270 - Norwegian
(+45) 4810-1120 - Spanish
(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications
Publications Customer Support
P.O. Box 29570
Raleigh, NC 27626-0570
USA

IBM Publications
144-4th Avenue, S.W.
Calgary, Alberta T2P 3N5
Canada

IBM Direct Services
Sortemosevej 21
DK-3450 Allerød
Denmark

- **Fax** — send orders to:

United States (toll free)
Canada
Outside North America

1-800-445-9269
1-403-267-4455
(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **On the World Wide Web**

Redbooks Web Site
IBM Direct Publications Catalog

<http://www.redbooks.ibm.com/>
<http://www.elink.ibm.com/pbl/pbl>

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number _____
- Credit card number _____

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

IBM	International Business Machines Corporation	KB	one thousand bytes
ITSO	International Technical Support Organization	TB	one billion bytes
PROFS	Professional Office System	HA	high availability
CISC	Complex Instruction Set Computing	HAS	high availability solution
RISC	Reduced Instruction Set Computing	ECS	electronic customer support
IMPI	Internal Microprogrammed Instruction	ERP	error recovery procedures
API	Application Programming Interface	PSP	preventive service planning
ADSM	ADSTAR Distributed Storage Manager	CPM	continuous power memory
CPU	central processing unit	CUM	cumulative PTF package
DOS	disk operating system	MB	one million bytes
IBM	International Business Machines	ISV	independent solution vendor
IT	information technology	APPC	advanced program to program communications
ITSO	International Technical Support Organization	RETAIN	REmote Technical Assistance and Information Network
LAN	local area network	URL	Uniform Resource Locator
IP	internal protocol	SRC	system reference code
TCP	transmission control protocol	APPN	application programming networking
		SNA	systems network architecture
		SNADS	SNA distribution services
		PTF	program temporary fix
		UPS	uninterruptible power supply

Index

Special Characters

/ file system (root) 223
/QFPNWSSTG directory 290
/QLanSrv directory 291
/QLanSrv objects 290
*AIX 64
*ALL 43
*BASE 64
*BLKSF (block special file object) 231
*BLKSF (block special file) 223
*CRITMSG action 150
*DSCENDRQS (disconnect end request) option 190
*DSCENDRQS action 150
*DSCMSG (disconnect message) option 190
*DSCMSG action 150
*ENDJOB action 150
*ENDJOBNO LIST action 151
*ENDSYS action 150
*ERR value 50
*FULL 43
*INTERACT 101
*MACHINE POOL 101
*MIN 43
*MINFIXLEN (minimum fixed length) option 332
*MSG action 150
*NETWARE 64
*PWRDWN SYS action 150
*REGFAC action 150
*RMVINTENT (remove internal entry) option 332
*SYS 43
#JOEVAT SMAPP task 39
#JOIJSS SMAPP task 39
#JOTUNT SMAPP task 39

Numerics

2683 adapter 352
2685 adapter 352
2726 MFIO P feature code 28
3590 tape device 30
64-bit addressing 33
9337 Disk Array Subsystem 28
9751 MFIO P 28
9751 MFIO P feature code 28

A

a-side of Licensed Internal Code (LIC) 168
abbreviations 421
abnormal IPL 33, 37, 47
access log 218
access log setup 219
access path journal entry 39

access path page images 332
access paths 38
 explicit journaling of 316
 journaling 316
acronyms 421
action-exit program 171
actions 162
activate LAN manager (ACTLANMGR)
 parameter 198
ACTLANMGR (activate LAN manager)
 parameter 198
Add Physical File Trigger (ADDPFTRG) command 306
additional value 144, 145
ADDPFTRG (Add Physical File Trigger) command 306
administrative client 110
administrative client function 112
ADSM (ADSTAR Distributed Storage Manager) 261,
 267
ADSM Notes backup agent 270
ADSM OS/2 Lotus Notes agent, saving with
ADSM/400 (ADSTAR Distributed Storage
 Manager/400) 110
 administrative client 110
 backup client 110, 111
 disaster recovery 112
 elements 110
 interoperability 114
 server 110
ADSTAR Distributed Storage Manager (ADSM) 261,
 267
ADSTAR Distributed Storage Manager/400
ADSTAR Distributed Storage Manager/400
 (ADSM/400) 110
alert filtering 136
allow virtual APPN support (ALWVRTAPPN)
 value 184
alternate installation device 30
alternate system 392
ALWVRTAPPN (allow virtual APPN support)
 value 184
Analyze Database File (ANZDBF) command 304
ANZDBF (Analyze Database File) command 304
API interface 353
APPC communications 2
APPC controller
 automatic creation 205
 automatic deletion 205
APPC controller description error recovery 204
APPC device description
 automatic creation 205
 automatic deletion 205
application client 351
Apply function 403

- apply journal changes 311
- apply process 399
- Apply PTF (APYPTF) command 177
- APPN High-Performance Routing (HPR) function 183
- APPN HPR (High-Performance Routing) 183
- APYPTF (Apply PTF) command 177
- archive function 111
- ARCserve for NetWare Version 6 285
- AS/400 (BRMS)
 - backup 114
 - media service 114
 - recovery 114
- AS/400 Forum
- AS/400 Model 400 system 387
- AS/400 NetServer 262
- AS/400 Performance Edge
- AS/400 Performance Management (PM/400)
- ASP
 - configuration 121
 - management 122
 - monitoring 121
- ASP (auxiliary storage pool) 56
- asynchronous delivery mode 331
- asynchronous replication 330
- audit journal 138
- Audit Journal Reader 404
- August 23, 1928 (system default date) 152
- authorization list 246
- auto configuration parameter 201
- autodelete function 205
- Automated Tools for System Management Functions
- automatic refresh interval parameter 139
- automatic tuner process 101
- automating message management 131
- automating security management 137
- Automation of Message Management
- Automation of Security Management
- auxiliary storage lower limit (QSTGLOWLMT) system value 149
- auxiliary storage lower limit action (QSTGLOWACN) system value 149
 - *CRITMSG action 150
 - *ENDSYS action 150
 - *MSG action 150
 - *PWRDWNSYS action 150
 - *REGFAC action 150
- auxiliary storage pool (ASP) 56
- availability
 - DASD 23
- availability options 63
- available value 145, 146

B

- b-side of Licensed Internal Code (LIC) 168
- BACKACC command 296
- backing up
 - specific objects from your Windows NT 260
 - system objects 259

- backing up (*continued*)
 - user objects 259
- backup
 - for AS/400 (BRMS) 114
 - integrated file system 221
 - internal battery 29
 - licensed program 85
 - PRPQ 85
 - server database 113
 - storage pool 113
 - strategy 61
- backup and recovery
 - plan 17
- backup and recovery plan 16
- backup client 110, 111
- backup domain controller (BDC) 263
- backup function 111
- backup function, directory-tree 111
- Backup Recovery and Media Services for AS/400
- backup to AS/400 tape 261
- BDC (backup domain controller) 263
- benchmark configuration 47
- bibliography 413
- block size 97
- block special file (*BLKSF) 223
- block special file object (*BLKSF) 231
- break handling program 132
- break handling program, creating 132
- BRM15A5 message 117
- BRM15A6 message 117
- BRMLOG report 117
- BRMS
 - maintenance activities 116
 - message handling 115
 - recovery report 115
 - skip-shipped 115
 - user exits 115
- BRMS/400 interoperability 114
- brown outs 28
- bus speed 99
- bus-level mirroring 25
- business impact analysis 14

C

- cannot allocate device message 187
- capacity limitations 361
- capture process 399
- cascading 330
- CCSID (character set ID) 72
- CCSID system 89
- Change Cleanup Options (CHGCLNUP) display 161
- change control server function 130
- Change IPL Attributes (CHGIPLA) command 34, 42
- change job (QWTCHGJB) API 165
- Change Journal (CHGJRN) command 313
- Change Message Description (CHGMSGD)
 - command 115

Change Message Queue (CHGMSGQ) command 132
Change Physical File (CHGPF) command 40
Change Shared Pool (CHGSHRPOOL) display 101
change shared storage pool (CHGSHRPOOL) API 166
character set ID (CCSID) 72
check job tables function 44
checksum 28
CHGASPDSC command 120
CHGIPLA (Change IPL Attributes) command 34, 42
CHGJRN (Change Journal) command 313
CHGMSGD (Change Message Description) command 115
CHGMSGQ (Change Message Queue) command 132
CHGPF (Change Physical File) command 40
CHGSHRPOOL (Change Shared Pool) display 101
CHGSHRPOOL (change shared storage pool) API 166
CHIGIPLA command 35
chip-to-chip circuitry test 35
CL command
 CHGASPDSC 120
 CPYLIBASP 120
 creating to save spooled files 70
 MOVLIBASP 120
 PRDBDP 120
 PRTASPLIB 120
 SAVLIBASP 120
 WRKASP 120
client access mode description 207
client characteristics 112
client failure 209
client-to-client optical link communication 356
client-to-server communication 356
clustered environments 327
clustering 351
co-requisite PTF 167, 175
cold start 37
command
 ADDPFTRG (Add Physical File Trigger) command 306
 ANZDBF (Analyze Database File) 304
 APYPTF (Apply PTF) 177
 BACKACC 296
 CHGIPLA (Change IPL Attributes) 34, 42
 CHGJRN (Change Journal) 313
 CHGMSGD (Change Message Description) 115
 CHGMSGQ (Change Message Queue) 132
 CHGPF (Change Physical File) 40
 CPYF (Copy File) 160
 CPYMEDIBRM 117
 CRTJRN (Create Journal) 313
 CRTNWSD (Create Network Server Description) 273
 CRTUDFS (Create User Defined File System) 234
 DSPCTLD (Display Controller Description) 183
 DSPJOBTL (Display Job Tables) 38, 143
 DSPMFSINF (Display Mounted File System Information) 233
 DSPMFSINF (Display Mounted FS Information) 235
 command (*continued*)
 DSPRCYAP (Display Recovery for Access Paths) 39, 41
 DSPRCYAP (Display Recovery for Access Paths) command 39
 DSPUDFS (Display User-Defined System) 236
 DUPTAP (Duplicate Tape) 80
 Edit Recovery for Access Paths (EDTRCYAP) 39
 EDTRCYAP (Edit Recovery for Access Paths) 39, 41
 EDTRCYAP (Edit Recovery for Access Paths) command 39
 ENDDOMSRV (End Domain Server) 251
 INSNTWSVR (Install NetWare Server) 272
 INSNWSAPP (Install Network Server Application) 268
 integrated file system 240
 LOOPBACK 214
 NETSTAT 214
 ObjectConnect/400 65
 PING 214
 PRTSYSINF (Print System Information)
 PWRDWNSYS (Power Down System) 43, 161
 QRYDOCLIB (Query Document Library) 241
 RCLDLO (Reclaim Document Library Object) 158
 RCLSPLSTG (Reclaim Spool Storage) 159
 RCLSTG (Reclaim Storage) 153
 RCVJRNE (Receive Journal Entry) 327
 RCVMSG (Receive Message) 132
 RESTACC 296
 RMVPTF (Remove PTF) 169
 RST (Restore) 223
 RSTDLO (Restore Document Library Object) 244, 247
 RSTLICPGM 86
 RTVCFGSRC (Retrieve Configuration Source) 268
 SAV (Save Object) 76
 SAV (Save) 223
 SAVCFG 54
 SAVCHGOBJ (Save Changed Objects) 76
 SAVDLO (Save Document Library Object) 247
 SAVDLO (Save Document Library Objects) 76
 Save Licensed Program (SAVLICPGM) 278
 SAVLIB (Save Library) 76
 SAVLICPGM 86
 SAVLICPGM (Save Licensed Program) 278
 SAVOBJ (Save Object) 76
 SAVSYS 53
 SAVSYSBRM 119
 SBMNWSCMD (Submit Network Server Command) 271
 SNDNETF (Send Network File) 177
 SNDPTFORD (Send PTF Order) 172
 STRMNTBRM 116
 TCP AS/400 alias CL 213
 VRYCFG (Vary Configuration) 186
 WRKACTJOB (Work With Active Jobs) 139
 WRKNTWVOL (Work NetWare Volume) 272

command (*continued*)
 WRKNTWVOL (Work With NetWare Volumes) 284
 WRKPCYBRM (Work With Move Policies) 117
 command line interface 271
 commitment control 51
 common format 219
 communications arbiter 182
 communications error recovery procedures
 (ERP) 181
 Communications Monitor 405
 communications recovery 203
 first-level recovery 203
 second-level recovery 203
 COMPACT (compaction) 383
 compaction (COMPACT) 383
 compaction algorithm 96
 compatibility
 DTACPR parameter 80
 TGTRLS parameter 79
 USEOPTBLK parameter 79, 80
 compress job tables 143
 compress job tables (CPRJOBTL) parameter 142
 compression (DTACPR) 383
 concurrent 55
 concurrent save 2, 50, 98, 246
 conditional PTF 175
 configuration ring services (CRS) function 198
 considerations, object locks 66
 continuous availability 17
 continuous operations 17
 continuous power main storage (CPM) 37
 continuously available 17
 continuously operational 17
 continuously powered main storage (CPM) 26
 control byte 97
 control panel 35
 control server 399
 controller level protection 24
 Copy File (CPYF) command 160
 copy server 399
 CPA2610 inquiry message 186
 CPA5316 message 133
 CPC2957 message 160
 CPC6260 message 161
 CPC8208 status message 156
 CPD0940 message 106
 CPD3244 message 160
 CPD3728 message 75
 CPD3754 message 81
 CPD376E message 81
 CPD378A message 79, 80
 CPD3796 message 60
 CPD6265 message 161
 CPF1187 message 187
 CPF1269 message 198
 CPF1273 message 187
 CPF1274 message 187
 CPF1275 message 187
 CPF2119 error message 157
 CPF2120 error message 157
 CPF2126 error message 157
 CPF2127 error message 157
 CPF2460 escape message 151
 CPF32A1 message 153
 CPF32A2 message 154
 CPF32A3 message 154
 CPF32A4 message 154
 CPF384E message 80
 CPF594C message 207
 CPF8113 message 153
 CPF8201 error message 157
 CPF8204 error message 157
 CPF8205 error message 157
 CPF8209 error message 157
 CPF8211 error message 157
 CPF8224 error message 157
 CPF8251 error message 157
 CPF8252 error message 157
 CPI099B message 150
 CPI099C message 150
 CPI1468 message 142
 CPI3712 message 51
 CPI3818 message 80
 CPI5970 message 207
 CPI8206 status message 156
 CPI8210 status message 156
 CPI8212 status message 156
 CPI8213 status message 156
 CPI8214 status message 156
 CPI8215 status message 156
 CPI8216 status message 156
 CPI8217 status message 156
 CPI8218 status message 156
 CPI8219 status message 156
 CPI8220 status message 156
 CPM (continuous power main storage) 37
 CPM (continuously powered main storage) 26, 29
 CPM function 29
 CPRJOBTL (compress job tables) parameter 142
 CPU usage
 on the source system 333
 on the target system 334
 CPYF (Copy File) command 160
 CPYLIBASP command 120
 CPYMEDIBRM command 117
 Create Journal (CRTJRN) command 313
 Create Network Server Description (CRTNWSD)
 command 273
 Create User Defined File System (CRTUDFS)
 command 234
 creating user spaces, considerations for backup when
 CRS (configuration ring services) function 198
 CRTJRN (Create Journal) command 313
 CRTNWSD (Create Network Server Description)
 command 273

CRTUDFS (Create User Defined File System)
 command 234
 cumulative (CUM) PTF package 168, 176
 cumulative backup, restoring changed objects
 from 257
 current-release-to-previous-release support 76

D

D-IPL 30
 DASD availability 23
 data compaction 95
 tape drive 95
 data compression (DTACPR) 58, 97
 data description specifications (DDS) format 219
 data replication 2, 397
 data replication process 399
 apply process 399
 capture process 399
 data server 399
 data transfer time 97
 database files, saving for recovery 303
 database journaling 310
 performance 312
 database protection 303
 database server 351
 database server job 322
 DataJoiner 398
 DataMirror Corporation 394
 High Availability Suite 394
 ObjectMirror 395
 OptiConnect 396
 SwitchOver System 395
 DataMirror HA Data 395
 high availability options 395
 DataPropagator 69
 configuration 399
 DataPropagator Relational Capture and Apply for
 AS/400 397
 DataPropagator Relational/400 393
 DataPropagator/400 397, 398
 DB2 Multisystem for OS/400 322
 DCDB (domain controller database) 292
 DDM (distributed data management) 351
 DDS (data description specifications) format 219
 Decision Control Matrix 395
 default on savxxx command 320
 delayed PTF 168
 device allocated message 187
 device I/O error action (QDEVRCYACN) system
 value 150
 *DSCENDRQS action 150
 *DSCMSG action 150
 *ENDJOB action 150
 *ENDJOBNO LIST action 151
 *MSG action 150
 device parity 23, 26, 27
 device parity protection 28

device recovery action (DEVRCYACN)
 parameter 190
 *DSCENDRQS (disconnect end request)option 190
 *DSCMSG (disconnect message) option 190
 DEVRCYACN (device recovery action)
 parameter 190
 DHCP (dynamic host configuration protocol) 212
 diagnostic
 hardware 35
 processor 35
 directory-tree backup function 111
 disaster recovery 112
 disconnect end request (*DSCENDRQS) option 190
 disconnect message (*DSCMSG) option 190
 disk failures 19
 disk space report 59
 Display Controller Description (DSPCTLD)
 command 183
 Display Device Description display 187
 Display Job Tables (DSPJOB TBL) command 38, 143
 Display Mounted File System Information
 (DSPMFSINF) command 233
 Display Mounted FS Information (DSPMFSINF)
 command 235
 Display Recovery for Access Paths (DSPRCYAP)
 command 39, 41
 Display User-Defined System (DSPUDFS)
 command 236
 distributed data management (DDM) 351
 distributed file
 backup example 324
 restoring 323
 distributed files, backup considerations 323
 distributed relational database architecture
 (DRDA) 401
 distributing PTFs 177
 distribution object, restoring 247
 Distribution Reader 404
 DLO
 maximum sequence number 50
 DLO (document library object) 241
 DLO(*SEARCH) parametric search 242
 DNS (domain name services) 212
 document library object (DLO) 241
 authority and ownership issues while
 restoring 247
 performance considerations 245
 restoring 244, 245
 saving 241
 document library services 241
 document library services file system (QDLS) 222
 domain controller
 restoring 263
 domain controller database (DCDB) 292
 saving 293
 domain name services (DNS) 212
 Domino database
 recovery examples 256

- Domino databases, limiting the location of 250
- Domino for AS/400 248
 - baking up 249
 - libraries and directories for 249
 - recovery 254
- Domino for AS/400 server
 - backing up 248, 250
 - backing up changed objects from 252
 - recovering the entire 254
 - restoring changed objects to 256
- Domino mail, recovering 255
- Domino server
 - backing up mail from 251
 - MAIL.BOX database 251
- Domino subdirectory, restoring changed objects to a specific 258
- DRDA (distributed relational database architecture) 401
- DSMNOTES program 271
- DSPCTLD (Display Controller Description) command 183
- DSPJOBTL (Display Job Tables) command 38, 143
- DSPMFSINF (Display Mounted File System Information) command 233
- DSPMFSINF (Display Mounted FS Information) command 235
- DSPRCYAP (Display Recovery for Access Paths) command 39, 41
- DTACPR (compression) 383
- DTACPR (data compression) 58, 97
- DTACPR (data compression) parameter compatibility 80
- dual path optical link 356
- Duplicate Tape (DUPTAP) command 80
- DUPTAP (Duplicate Tape) command 80
- dynamic host configuration protocol (DHCP) 212
- dynamic priority scheduling 105

E

- ECS (electronic customer support) 168
 - PTF delivery 178
- Edit Rebuild of Access Paths (EDTRBDAP) display 38
- Edit Recovery for Access Paths (EDTRCYAP) command 39, 41
- EDTRBDAP (Edit Rebuild of Access Paths) display 38
- EDTRCYAP (Edit Recovery for Access Paths) command 39, 41
- electronic customer support (ECS) 168
- emergency power off 37
- End Domain Server (ENDDOMSRV) command 251
- End Job Abnormal (ENDJOBABN) command 152
- end subsystem option (ENDSBSOPT) parameter 162
- ENDDOMSRV (End Domain Server) command 251
- ENDJOBABN (End Job Abnormal) 152
- ENDJOBABN (End Job Abnormal) command 152
- ENDSBSOPT (end subsystem option) 162
- ENDSBSOPT (end subsystem option) parameter 162

- Entries field 145
- ERP (error recovery procedures) 181
- error log 218
- error log filtering 199
- error log setup 220
- error logging
 - protocol errors 189
- error recovery
 - TCP 188
 - testing 208
- error recovery failure
 - client failure 209
 - network failure 208
 - server failure 209
 - types 208
- error recovery procedures (ERP) 181
 - testing tips 209
- error, human 19
- exit program 327
- expert cache function 104
- explicit journaling of access paths 316
- extended main storage diagnostic test 35

F

- failure
 - disk 19
 - program 19
 - system 19
- faults-per-second parameter 103
- fiber bus hardware 351
- file level restore 263
- file server file system (QFileSvr.400) 223
- file server job restructure 197
- filter rule
 - saving and restoring with the COPY command 302
- firewall 297
 - creating a library for back-up files 299
 - restoring 300
 - restoring communication configuration objects 301
 - restoring operational data 301
 - restoring the configuration data 302
 - saving 298
 - saving communication configuration objects 299
 - saving the configuration 300
 - saving the operational data 300
 - stopping 299
- firewall NWSD
 - varying off 299
 - varying on 300
- folder (FLR) parameter 56
- force vary off (FRCVRYOFF) parameter 186
- force write ratio 40
- FRCVRYOFF (force vary off) parameter 186
- full programmed 36

G

global log setup 219
graphical workstation interface 271

H

HA data, DataMirror 395
hang, CPM or main store dump processing 152
hardware availability options 23
hardware data compression (HDC) 95
 DTACPR 96
 HDC algorithm 96
hardware diagnostics 35
hardware diagnostics (HDWDIAG) parameter
 *ALL 43
 *MIN 43
hardware disk compression 73
hardware diagnostics (HDWDIAG) parameter 43
HAS/400 (High Availability Services) 409
HDC (hardware data compression) 95
HDC algorithm 96
HDWDIAG (hardware diagnostics) parameter 43
Hierarchical Storage Management (HSM) 73
high availability 17, 24, 75, 329
 features 393
 functions 393
 RAID-V 27
high availability application 2
high availability clustered solution 351
High Availability Services (HAS/400) 409
High Availability Solution 69
 consideration 392
 DataMirror 393
 DataPropagator Relational Capture and Apply for AS/400 397
 example 391
 Lakeview Technology 393
 Vision Solutions, Inc. 393
High Availability Solutions
high impact or pervasive (HIPER) PTF 168
high systems availability 391
High-Performance Routing (HPR) 183
high-speed optical bus 354
highly available 17
HIPER (high impact or pervasive) PTF 168, 169
hog hunter function 105
hot site backup 69
HPR (High-Performance Routing) 183
HSM (Hierarchical Storage Management) 73
HTML object authority 217
HTTP files, backing up 217
HTTP log file setup 219
HTTP server configuration 217
HTTP server protection setup 217
hub 351
hub dual path connection 357
hub system 387

hub-to-hub communication 357

I

IBM DataJoiner 398
IBM Job Scheduler for OS/400
IBM System View System Manager for AS/400
IBM SystemView Managed System Services for AS/400
IDRC drive 96
immediate PTF 168
In-use entries field 145
 available value 146
 in-use value 146
 other value 146
 size value 145
 table value 145
 total value 146
in-use value 146
incremental backup, restoring Domino database from 257
incremental function 111
incremental saves 61
initial value 144, 145
INSNTWSVR (Install NetWare Server) command 272
INSNTWSVR method 282
INSNWSAPP (Install Network Server Application) command 268
Install NetWare Server (INSNTWSVR) command 272
Install Network Server Application (INSNWSAPP) command 268
integrated file system 221
 commands 240
 restore 223
 save 223
 structure 221
integrated file system object, saving 212
Integrated PC Server
 saving a directory for 293
Integration for NetWare program
 restoring 284
 saving 278
internal battery backup 29
internal journal entry, remove 332
internal journal receivers 39
Internet connection servers 213
Internetwork Packet Exchange (IPX) support
 saving and restoring 286
interoperability
 ADSM/400 114
 BRMS/400 114
IOP level 25
IPL 30, 33
 abnormal 33, 37, 47
 affecting the time to 36
 benchmarks 46
 changing attributes 42
 cleanup tasks 33
 fast (default) 35

IPL (*continued*)
 full programmed 36
 hardware configuration 37
 manual 37
 marking progress with SRC codes 44
 modes 34
 normal 37, 371
 performance 33
 performance improvements 44
 process 34
 programmed 37
 short programmed 35
 slow 35
 software configuration 37
 startup tasks 33
 types 34
 IPL time 371
 IPX circuit entry, saving 277

J

job message queue full action (QJOBMSGQFL) system value 151
 job scheduler (OS/400) 125
 Job Scheduler/400 148
 job scheduling 127
 journal entry
 asynchronous replication 330
 synchronous replication 330
 journal entry latency 331
 journal receiver ASP, create 331
 journal receiver protection 321
 journal receivers 83
 journal replication 327
 journaling of access paths
 journaling, performance tips 313

L

Lakeview Technology 401
 MIMIX/400 401, 402
 MIMIX/Monitor 401
 MIMIX/Object 401
 MIMIX/Promoter 401
 MIMIX/Switch 401
 LAN administrator authority 291
 LAN response timer (LANRPTMR) parameter 186
 LAN Server for OS/400 (OS/2 Warp Server for AS/400) 287
 LAN Server/400, restoring 296
 language, secondary 88
 LANRPTMR (LAN response timer) parameter 186
 level of service 17
 levels of protection 18
 library file system (QSYS.LIB) 222
 library in use message 53
 library names 2, 89
 list of 85

LIC (Licensed Internal Code) 23, 30
 LIC PTF apply 171
 LIC trace 191
 Licensed Internal Code (LIC) 23, 30
 a-side 168
 b-side 168
 licensed program
 backup 85
 considerations 85
 installation considerations 88
 library names 89
 recovery 85
 restore considerations 88
 restore methods 87
 save methods 86
 saving 278
 user profile authority 89
 licensed program library
 restoring commands to QSYS 92
 licensed program product (LPP) 85
 restore 85
 save 85
 LICPGM menu 87
 link 387
 link redundancy 388
 list job (QUSLJOB) API 165
 load source 24
 load source mirroring
 protection 25
 remote 25
 load source unit 23
 lock conflict
 CHGJRN 315
 RCVJRNE 315
 locking conflicts 2
 locks considerations, object 66
 logical dependencies 51
 logical file 304
 Logical Switch 405
 LOOPBACK command 214
 Lotus Notes
 backup 267
 recovery 267
 Lotus Notes on the Integrated PC Server 264
 LPP (licensed program product) 85
 LZ1 compression algorithm 73

M

mail, restoring 247
 main store dump 47
 manual IPL 37
 MAXFRAME (maximum frame size) value 185
 maximum frame size (MAXFRAME) value 185
 maximum tape performance 98
 media class 117
 CPYMEDIBRM command 117
 media errors 53

media order 179
media service 114
message
 BRM15A5 117
 BRM15A6 117
 cannot allocate device 187
 CP13712 51
 CPA2610 186
 CPA5316 133
 CPC2957 160
 CPC6260 161
 CPD0940 106
 CPD3244 160
 CPD3728 75
 CPD3754 81
 CPD376E 81
 CPD378A 79, 80
 CPD3796 60
 CPD6265 message 161
 CPF1187 187
 CPF1269 198
 CPF1273 187
 CPF1274 187
 CPF1275 187
 CPF2460 151
 CPF32A1 153
 CPF32A2 154
 CPF32A3 154
 CPF32A4 154
 CPF384E 80
 CPF595C 207
 CPF8113 153
 CPI099B 150
 CPI099C 150
 CPI1468 142
 CPI3818 80
 CPI5970 207
 device allocated 187
 library in use 53
 object not found 89
 RCLSTG (reclaim storage) status 156
 reclaim storage error 157
MFIOP (multi-function I/O processor) 24
 initialization 36
microcode PTF 170
MIMIX with OptiConnect 406
MIMIX/400 401, 402
 Apply function 403
 Receive function 403
 Send function 402
 Switch function 403
 Synchronize function 403
MIMIX/Monitor 401, 405
 prepackaged monitor program 406
MIMIX/Object 401, 404
 Audit Journal Reader 404
 Distribution Reader 404
 Send Network Object 404
 MIMIX/Promoter 401, 406
 MIMIX/Switch 401, 405
 Communications Monitor 405
 Logical Switch 405
 Physical Switch 405
 mini-reclaim storage 155
 minimizing save and restore time 379
 minimum fixed length (*MINFIXLEN) option 332
 minimum switched status (MINSWTSTS) 202
 MINSWTSTS (minimum switched status) 202
 mirrored protection 24
 bus level 25
 remote 25
 mirroring 23
 bus-level 25
 Model 400 system 387
 mounted UDFS 237
 restoring 240
 saving 237
 MOVLIBASP command 120
 MSS/400 (SystemView Managed System Services for AS/400) 130
 multi-function I/O processor (MFIOP) 24
 multi-member database file save performance 321
 multilingual support 72
 multilingual system environment 72
 multinational environment 325
 multinational system 89
 multiple documents, saving 242
 multiple file systems
 restoring across 228
 saving across 224

N

NDS** (Netware Directory Services**) 272
NETSTAT command 214
NetWare configuration
 restoring 282
 saving 275
NetWare data
 restoring 274
 restoring from QNetWare directory 285
 saving 274
 saving from the QNetWare directory 280
Netware Directory Services** (NDS**) 272
NetWare file system (QNetWare) 223
NetWare on Integrated PC Server 272
NetWare server configuration
 saving 278
NetWare volume 279
 restoring from /QFPNWSSTG directory 284
 saving from /QFPNWSSTG directory 279
network server
 storage space 272
 volume 272
network server description
 restoring 269
 saving 268

- network server storage space 274, 290
 - restoring 270
 - saving 269
- network storage space
 - restoring 284
 - saving 279, 294
- network time zone synchronization (BRMS) 117
- nightly backup, restoring changed objects from 257
- no lock option 51
- non-updateable action PTF 171
- normal IPL 37, 371
- Notes agent, restoring databases 271
- Notes backup agent
 - backing up 271
 - DSMNOTES program 271
 - user interface 271

O

- Object Distribution System (ODS/400) 408
- object in use message 53
- object locks considerations 66
- Object Mirroring System (OMS/400) 407
- object not found message 89
- ObjectConnect/400 65
 - availability 65
 - benefits 68
 - command sets 66
 - implementation considerations 69
- ObjectConnect/400 command 65
- ObjectMirror 395
- observability 77
 - removing 78
- observable object 77
 - program template 77
- ODS/400 (Object Distribution System) 408
- office services information, saving 243
- omit options 2
- OMIT parameter 119
 - SAVSYSBRM command support for 119
- OMS/400 (Object Mirroring System) 407
- online at IPL parameter 201
- open systems file system (QOpenSys) 222
- Operational Assistant
 - Operational Assistant menu 134
 - operational LAN manager activation 198
- Operations Control Center/400 128
- optical bus cable 100
- optical file system (QOPT) 223
- optical link hardware 353
- OptiConnect 396
 - dual path connection 357
 - hub selection 356
 - link 387
 - path 387
 - RPQ 359
 - satellite 355
 - satellite dual path connection 356
 - single path 356

- OptiConnect (*continued*)
 - with MIMIX 406
- OptiConnect cluster 352, 387
- OptiConnect for OS/400 351
 - remote journal 353
- OptiConnect hardware 352
- OptiConnect terminology 387
- opticonnect/400 67
- OptiMover 353
- orphan 154
- orphan data 21
- OS/2 Warp server
 - saving specific objects 293
- OS/2 Warp Server for AS/400
 - backup and recovery 290
 - restoring 295
 - saving on RISC machines 291
- OS/2 Warp Server for AS/400 (LAN Server for OS/400)
 - saving 287
 - saving objects with multiple names 292
- OS/2 Warp Server for AS/400 file system (QLanSrv) 222
- OS/2 Warp Server for AS/400 structure 288
- OS/400 Alert Support 135
- OS/400 Job Scheduler 125
- OS/400 Security Tools 137
- other value 146
- outages 17

P

- Packet Internet Groper (PING) 214
- parameter
 - ACTLANMGR (activate LAN manager) 198
 - auto configuration 201
 - automatic refresh interval 139
 - CPRJOBTL (compress job tables) 142
 - DEVRCYACN (device recovery action) 190
 - faults-per-second 103
 - folder (FLR) 56
 - FRCVRYOFF (force varyoff) 186
 - LANRSPTMR (LAN response timer) 186
 - OMIT 53, 119
 - online at IPL 201
 - priority 102
 - PRTVSNRPT 116
 - RCVSIZOPT (receiver size options) 332
 - recovery locations 116
 - RSTFLR (restore folder) 247
 - RTVVOLSTAT 116
 - RUNCLNUP 116
 - SAVACT (save while active) 75
 - size percentage 102
 - TGTRLS (target release) 75
 - TIMOUTOPT (timeout option) 162
 - use optimum block size (USEOPTBLK) 97
 - USEOPTBLK (use optimum block) 57, 78
- parametric search 242

- parity protection 28
- pass-through server job 195
- path 387
- path redundancy 388
- PC client 297
- PDC (primary domain controller) 263
- performance management 18
- Permanent job structures field 144
 - additional value 144
 - initial value 144
- permanent PTF 169
- physical file 304
- physical link 352
- Physical Switch 405
- PING (Packet Internet Groper) 214
- PING command 214
- point of failure 21
- Power Down System (PWRDWN SYS) command 43, 161
- power down times 29
- PRDBDP command 120
- prepackaged monitor program (MIMIX/Monitor) 406
- prerequisite PTF 167, 175
- prestart jobs 206
- preventive service planning (PSP) 168, 169
- previous release system 83
- primary bus receiver 354
- primary domain controller (PDC) 263
- Print System Information (PRTSYSINF) command 123
- print system information parameter 65
- priority parameter 102
- priority, sending task 334
- processor diagnostics 35
- product-level PTF 171
- production system 392
- program start request message threshold 198
- program template 77
- program temporary fix (PTF) 167
- programmed IPL 37
- protection
 - controller level 24
 - device parity 28
 - mirrored 24
 - remote mirrored 25
- protocol error logging 189
- PRPQ
 - backup 85
 - considerations 85
 - recovery 85
 - WRKASP utility for OS/400 120
- PRTASPLIB command 120
- PRTSYSINF (Print System Information) command 123
- PRTVSNRPT parameter 116
- PSP (preventive service planning) 168, 169
- PTF 54
- PTF (program temporary fix) 167
- PTF cover letter 167

- PTF cover letter order 179
- PTF package 2
- PTFs for CHGJRN performance improvement 314
- PWRDWN SYS (Power Down System) command 43, 161
- PWRDWN SYS RESTART (*YES *YES) command 35

Q

- QAUTOVRT system value 188
- QCMNARB system value 182
- QDEVRCYACN (device I/O error action) system value 150
- QDFTJOBSCD job-schedule object 148
- QDLS (document library services file system) 222
- QDLS file system, restoring objects to 230
- QDLS physical file system 241
- QFileSvr.400 (file server file system) 223
- QFPINT library
 - QFBSYS4 AS/400 object name 274
 - QFPBSYS2 AS/400 object name 273
 - QFPBSYS2 storage space 264
- QGPL library 81
- qjoaddressremotejournal API 334
- QJOBMSGQFL (job message queue full action) system value 151
- QJOBSCD job scheduler 148
- qjochangerejournalstate API 334
- qjoremovejournalstate API 334
- qjoretrievejournalinformation API 334
- qjortvjrnreceiverinformation API 334
- QLanSrv (OS/2 Warp Server for AS/400 file system) 222
- QlpHandleCdState API 80
- QNetWare (NetWare file system) 223
- QNetWare characteristics 272
- QOpenSys (open systems file system) 222
- QOPT (optical file system) 223
- QPASTHRSVR system value 194
- QPFRAJ system job 100
- QPSRVDMP spooled file 371
- QRYDOCLIB (Query Document Library) command 241
- QSRSAVO API 78
 - USEOPTBLK parameter (V3R7) 78
- QSTGLOWACN (auxiliary storage lower limit action) system value 149
- QSTGLOWLMT (auxiliary storage lower limit) system value 149
- QSYS physical file 159
- QSYS.LIB (library file system) 222
- QSYS.LIB file system
 - restoring objects from 229
 - saving objects from 226
- QSYSARB system arbiter job name 147
- QSYSARBN system arbiter job name 147
- Qsysinc implementation 335
- QSYSMSG message queue 131
 - creating 131

QSYSOPR message queue 151
 removing obsolete messages 188
 Query Document Library (QRYDOCLIB)
 command 241
 QUSCHGPA Change Pool Attributes API 166
 QUSLJOB (list job) API 165
 QUSRJOB (retrieve job information) API 164
 QUSRSYS library 81, 274
 QYNALNCD storage space 265
 QYNASYS1 storage space 265
 SERVER 13 storage space 265
 SERVER11 storage space 264
 QWCBTCLNUP (work control block table
 cleanup) 143
 QWCCRTEC tool 378
 QWCCTREC tool 46
 QWTCHGJB (change job) API 165

R

RAID support 28
 RAID-V 27
 RCLDLO (Reclaim Document Library Object)
 command 158
 RCLSPLSTG (Reclaim Spool Storage) command 159
 RCLSTG (Reclaim Storage) command 153, 156
 CPC8208 status message 156
 CPI8206 status message 156
 CPI8210 status message 156
 CPI8212 status message 156
 CPI8213 status message 156
 CPI8214 status message 156
 CPI8215 status message 156
 CPI8216 status message 156
 CPI8217 status message 156
 CPI8218 status message 156
 CPI8219 status message 156
 CPI8220 status message 156
 RCLSTG (reclaim storage) status message 156
 RCVJRNE (Receive Journal Entry) command 327
 RCVMSG (Receive Message) command 132
 RCVSIZOPT (receiver size options) parameter 332
 Receive function 403
 Receive Journal Entry (RCVJRNE) command 327
 Receive Message (RCVMSG) command 132
 receiver size options (RCVSIZOPT) parameter 332
 Reclaim Document Library Object (RCLDLO)
 command 158
 Reclaim Spool Storage (RCLSPLSTG) command 159
 Reclaim Storage (RCLSTG) command 153
 reclaim storage error message 157
 CPF2119 157
 CPF2120 157
 CPF2126 157
 CPF2127 157
 CPF8201 157
 CPF8204 157
 CPF8205 157
 CPF8209 157

reclaim storage error message (*continued*)
 CPF8211 157
 CPF8224 157
 CPF8251 157
 CPF8252 157
 recommendation 14, 28, 35, 40, 44, 54, 73, 76, 89, 96,
 105, 118, 152, 163, 182, 189, 197, 201, 202, 205, 206,
 209, 224, 268, 302, 315, 319, 320, 331
 recovery
 for AS/400 (BRMS) 114
 licensed program 85
 PRPQ 85
 server database 113
 storage pool 113
 testing 61
 recovery locations parameter 116
 recovery plan 14
 recovery steps 20
 recovery time 19
 reducing the length of 332
 redundant fiber link solutions 352
 redundant optical link 356
 referential integrity, save and restore
 considerations 306
 relational database directory
 entries 330
 saving and restoring 307
 release-to-release support 75
 REM (ring-error monitor) function 198
 remote journal 353
 remote journal function 2, 327, 391, 396
 remote journal implementation 331
 remote journal replication mode 330
 remote journal transport protocol 330
 remote journals
 implementing with C 335
 implementing with RPG 339
 remote load source 23
 remove internal entry (*RMVINTENT) option 332
 Remove PTF (RMVPTF) command 169
 removing obsolete messages in QSYSOPR 188
 replication, data 397
 RESTACC command 296
 RESTART (restart type) parameter 43
 restart type (RESTART) parameter 43
 *FULL 43
 *SYS 43
 type 36
 restore 49
 commands from licensed program library to
 QSYS 92
 concurrent operations 57
 integrated file system 220
 licensed program 87
 licensed program product (LPP) 85
 observability 77
 observable object 77
 performance 95, 98

- restore (*continued*)
 - spooled files 69
- Restore (RST) command 223
- Restore Document Library Object (RSTDLO)
 - command 244, 247
- restore folder (RSTFLR) parameter 247
- restoring an individual document 272
- restoring NetWare server configuration 284
- restoring NWSD 282
- restoring server storage space 284
- restricted state 291
- restrictions 361
- RETAIN 167, 178
- Retrieve Configuration Source (RTVCFGSRC)
 - command 268
- retrieve job information (QUSRJOBI) API 164
- RETRIEVE method 282
- ring-error monitor (REM) function 198
- RMVPTF (Remove PTF) command 169
- root (/ file system) 223
- RPQ 85
- RST (Restore) command 223
- RSTDLO (Restore Document Library Object)
 - command 244, 245, 247
 - authority for 246
- RSTFLR (restore folder) parameter 247
- RSTLICPGM command 86
- RTVCFGSRC (Retrieve Configuration Source)
 - command 268
- RTVVOLSTAT parameter 116
- RUNCLNUP parameter 116

S

- SAM/400 (System Availability Monitor) 408
- satellite dual path connection 356
- satellite system 351, 387
- SAV (Save Object) command 76
- SAV (Save) command 223
 - USEOPTBLK parameter (V3R7) 78
- SAVACT (Save While Active) option 53
- SAVACT (save while active) parameter 75
- SAVCFG command 54
- SAVCFG(Save Configuration) command
 - USEOPTBLK parameter (V4R1) 78
- SAVCHGOBG (Save Changed Objects) command
 - USEOPTBLK parameter (V3R7) 78
- SAVCHGOBJ (Save Changed Objects) command 76
- SAVDLO (Save Document Library Object)
 - command 247
 - USEOPTBLK parameter (V4R1) 78
- SAVDLO (Save Document Library Objects)
 - command 76
- SAVDLO command, authority for 243
- save 49
 - authority requirements 290
 - concurrent 50, 98
 - concurrent for DLOs 56
 - concurrent operations on libraries 55

- save (*continued*)
 - incremental 61
 - licensed program 86
 - licensed program product (LPP) 85
 - performance 78, 95, 98
 - postprocessing 55
 - preprocessing 55
 - spooled files 69
 - start time 63
 - strategy 59
 - unattended 61
 - USEOPTBLK parameter 78
 - version-based 111
- Save (SAV) command 223
- save and restore rates 379, 380, 383
- save and restore spooled file (SAVRSTSPLF) tool 69
- Save Changed Objects (SAVCHGOBJ) command 76
- Save Document Library Object (SAVDLO)
 - command 247
- Save Document Library Objects (SAVDLO)
 - command 76
- Save Library (SAVLIB) command 76
- Save Licensed Program (SAVLICPGM) command 278
- save list 55
- SAVE menu 65, 85
 - unattended saves 61
- Save Object (SAV) command 76
- Save Object (SAVOBJ) command 76
- Save While Active (SAVACT) function 76
- Save While Active (SAVACT) option 53
- save while active (SAVACT) parameter 75
- Save While Active (SWA) function 2, 50, 77
 - overview 52
 - restrictions 52
 - shadow file 52
 - side file 52
- Save While Active function
 - checkpoint 51
 - checkpoint processing 51
- saving access path 318
- saving licensed program 278
- saving mail 244
- SAVLIB (Save Library) command 76
 - USEOPTBLK parameter (V3R7) 78
- SAVLIBASP command 120
- SAVLICPGM (Save Licensed Program) command 278
- SAVLICPGM command 86
- SAVOBJ (Save Object) command 76
 - USEOPTBLK parameter (V3R7) 78
- SAVRSTSPLF (save and restore spooled file) tool 69
- ZRSTSPLF command 69
- ZSAVSPLF command 69
- SAVSAVFDTA(Save Save File Data) command
 - USEOPTBLK parameter (V4R1) 78
- SAVSECDA (Save Security Data) command
 - USEOPTBLK parameter (V4R1) 78
- SAVSYS (Save System) command
 - USEOPTBLK parameter (V4R1) 78

SAVSYS command
 OMIT parameter 53
 omitting objects on 53
SAVSYSBRM command 119
SBACKUP restore option 285
SBMNWSCMD (Submit Network Server Command)
 command 271
scatter loading 23
SDC (software data compression) 96
search index database 246
secondary language 88
selective function 111
Send function 402
Send Network File (SNDNETF) command 177
Send Network Object 404
Send PTF Order (SNDPTFORD) command 172
sending task priority 334
server 110
server database, backup and recovery 113
server function 112
server message block (SMB) 262
server storage space 273, 290
 C: drive 264, 273
 D: drive 264, 273
 E: drive 264, 273
 F: drive 265, 274
 G: drive 265
 K: drive 265
 restoring 269
 saving 268, 278
 types 264
service reference code (SRC) 106
service time 98
settings for optimum performance 98
shadow file 52
SHRRD option 51
side file 52
single-level storage 23
site loss 19
size percentage parameter 102
size value 145
sizing tape requirements 381
skip-shipped 115
sleep mode 30
SLIC (System Licensed Internal Code) 30
SMAPP
 *MIN 48
 #JOEVAT task 39
 #JOIJSS task 39
 #JOTUNT task 39
 modifying 40
 performance considerations 40
 tasks 39
SMAPP (system managed access path
 protection) 38, 316
SMB (server message block) 262
SNADS (systems network architecture distribution
 services) 38
SNDNETF (Send Network File) 177
SNDPTFORD (Send PTF Order) command 172
software data compression (SDC) 96
 DTACPR 96
solution
 High Availability Solutions 391
 redundant fiber link 352
 solution, high availability clustered 351
 source sink component trace 191
 SPCN (System Power Control Network) 29
 Specify Command Defaults display 62
 SRC (service reference code) 106
 SRC codes 44, 371
 SST (system service tools) LIC interface 191
 steps, recovery 20
 storage pool
 backup and recovery 113
 storage space 290
 network server storage space 274
 server storage space 273
 types 273
 stored procedures 310
 STRMNTBRM command 116
 Submit Network Server Command (SBMNWSCMD)
 command 271
 subsystem configuration 200
 superseded PTF 167
 SWA (Save While Active) function 50
 Switch function 403
 switched disconnect 202
 SwitchOver System 395
 Decision Control Matrix 395
 Synchronize function 403
 synchronous delivery mode 330
 synchronous replication 330
 system availability
 work management 139
 System Availability Monitor (SAM/400) 408
 system checksum 28
 system crash 318
 system date 152
 system default date (August 23, 1928) 152
 system failure 19
 System Licensed Internal Code (SLIC) 30
 system managed access path protection
 (SMAPP) 38, 316
 system management, tools for automating 109
 System Power Control Network (SPCN) feature 29
 system recovery 318
 system redundancy 327
 system reply list 134
 system service tools (SST) LIC interface 191
 system tuner, tuning 104
 systems network architecture distribution services
 (SNADS) 38
 SystemView Managed System Services for AS/400
 (MSS/400) 130

T

table value 145
 tape devices 30
 tape drive
 data compaction 95
 ratings 96
 tape file sequence number 119
 tape requirements, sizing 381
 target 5250 display station pass through 191
 target release (TGTRLS) parameter 75
 target work station function 192
 task priority, sending 334
 task scheduler 105
 TCP AS/400 alias CL commands 213
 TCP error recovery 188
 TCP function, logical file 211
 TCP/IP common problems 215
 TCP/IP configuration
 saving 211, 299
 TCP/IP network stability 211
 TCP/IP tips 213
 Telnet devices 189
 device recovery action 190
 inactivity timeout 189
 Temporary job structures field 145
 additional value 145
 available value 145
 initial value 145
 temporary PTF 169
 test
 chip-to-chip 35
 extended main storage diagnostic 35
 testing recovery procedures 61
 text search service, recovery of text index files
 for 248
 text search services files, saving 244
 TGTRLS (target release) parameter 75
 compatibility 79
 The PRSYSINF Command
 The SystemView Launch Window for OS/400
 time to save 2
 timeout option (TIMOUTOPT) parameter 162
 TIMOUTOPT (timeout option) parameter 162
 actions 163
 total value 146
 transfer time of data 97
 transport mechanisms 330
 transport method 333
 trigger program, save and restore tips 306

U

UDFS (user-defined file system) 223, 231
 UDFS (user-defined file systems) 61
 unavailability 167

uninterruptible power supply (UPS) 28
 unmount file system 64
 unmounted UDFS 235
 restoring 237
 restoring an individual object from 237
 saving 236
 updateable action PTF 171
 UPS (uninterruptible power supply) 28
 use optimum block (USEOPTBLK) parameter 57, 78
 use optimum block size (USEOPTBLK) parameter 97
 USEOPTBLK (use optimum block size) parameter 97
 USEOPTBLK (use optimum block) parameter 57, 78
 compatibility 79
 USEOPTBLK(use optimum block) parameter
 compatibility 80
 considerations 80
 user class parameter 89
 user data 261
 user profile authority 89
 user storage space 263
 user-defined file system (UDFS) 223, 231
 mounting 231
 user-defined file systems (UDFS) 61

V

Vary Configuration (VRYCFG) command 186
 vary off network server parameter 63
 *ALL 64
 *LANSERVER 64
 *NONE 64
 vary off the network server parameter
 Version 1 (V1R1, V1R2, and V1R3) 2
 Version 2 (V2R1, V2R2, V2R3, and V3R05) 2
 Version 3 (V3R1, V3R2, V3R6, and V3R7) 2
 Version 4 (V4R1 and V4R2) 2
 version-based save 111
 virtual controller autcreate device (VRTAUTODEV)
 value 184
 Vision Solutions, Inc. 406
 High Availability Services (HAS/400) 409
 Object Distribution System (ODS/400) 408
 Object Mirroring System (OMS/400) 407
 System Availability Monitor (SAM/400) 408
 Vision Suite 407
 Vision Suite 407
 VLIC log 106
 VLOG improvement 190
 VLOG log 106
 volume media management 114
 VRTAUTODEV (virtual controller autcreate device)
 value 184
 VRYCFG (Vary Configuration) command 186

W

WCBT (work control block table) 141
 WCBTE (work control block table entry) 142

- white outs 29
- Windows NT
 - directories and objects for 258
 - operating system and registry 263
 - restoring 262
 - saving and restoring 258
- Windows NT operating system and registry 260
- Windows NT server
 - re-install 263
 - system files 262
- work control block table (WCBT) 141
- work control block table cleanup (QWCBTCLNUP) 143
- work control block table entry 143
- work control block table entry (WCBTE) 142
- work management 139
- work management API 164
- Work NetWare Volume (WRKNTWVOL) command 272
- Work With Active Jobs (WRKACTJOB) command 139
- Work With Job Schedule Entries display 125
- Work With Move Policies (WRKPCYBRM) command 117
- Work With NetWare Volumes (WRKNTWVOL) command 284
- Work With Shared Pools (WRKSHRPOOL) display 101
- write cache 331
- write capability 99
- WRKACTJOB (Work With Active Jobs) command 139
- WRKASP command 120
- WRKASP Utility for OS/400
- WRKASP utility for OS/400 PRPQ 120
- WRKASP utility, CL commands 120
- WRKCFGSTS display change 194
- WRKNTWVOL (Work NetWare Volume) command 272
- WRKNTWVOL (Work With NetWare Volumes) command 284
- WRKPCYBRM (Work With Move Policies) command 117
- WRKSHRPOOL (Work With Shared Pools) display 101

ITSO Redbook Evaluation

The System Administrator's Companion to AS/400 Availability and Recovery
SG24-2161-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

Customer **Business Partner** **Solution Developer** **IBM employee**
 None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

